working paper
department
of economics

PATHOLOGICAL MODELS OF OBSERVATIONAL LEARNING

Lones Smith
Peter Sørensen

94-24 July 1994

massachusetts
institute of
technology

50 memorial drive
cambridge, mass. 02139

PATHOLOGICAL MODELS
OF OBSERVATIONAL LEARNING

Lones Smith
Peter Sørensen

94-24 **July 1994**

Pathological Models of Observational Learning*

Lones Smith†
and
Peter Sørensen‡

Department of Economics
Massachusetts Institute of Technology
Cambridge, MA 02139

(Preliminary version. Comments encouraged.)

July 27, 1994

Abstract

This paper systematically analyzes the observational learning paradigm of Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992). We first relax the informational assumption that is the linchpin of the 'herding' results, namely that individuals' private signals are uniformly bounded in their strength. We then investigate the model with heterogeneous preferences, and discover that a 'twin' observational pathology generically appears: Optimal learning may well lead to a situation where no one can draw any inference at all from history! We also point out that counterintuitively, even with a constant background "noise" induced by trembling or crazy individuals, public beliefs generically converge to the true state of the world.

All results are cast within a simple dynamic mathematical framework that is (i) rich enough to describe a rich array of observational learning dynamics; and (ii) amenable to economic modifications that hinder or abet informational transmission, and sometimes permit full belief convergence to occur.

*The authors wish to thank participants at the MIT theory lunch, in particular Abhijit Banerjee, for their comments on a somewhat related paper. We are especially grateful to Chris Avery, who was very instrumental at an early stage in this project. All errors remain our responsibility. Sørensen gratefully acknowledges financial support for this work from the Danish Social Sciences Research Council.

†e-mail address: lones@lones.mit.edu

‡e-mail address: petersor@mit.edu

1. INTRODUCTION

Suppose that a countable number of individuals each must make a once-in-a-lifetime binary decision[1] — encumbered solely by uncertainty about the state of the world. Assume that preferences are identical, and that there are no congestion effects or network externalities from acting alike. Then in a world of complete and symmetric information, all would ideally wish to make the same decision.

But life is more complicated than that. Assume instead that the individuals must decide sequentially, all in some preordained order. Suppose that each individual may condition his decision both on his (endowed) private signal about the state of the world and on all his predecessors' decisions, but not their private signals. The above simple framework was independently introduced in Banerjee (1992) and Bikhchandani et al. (1992) (hereafter denoted BHW). Their perhaps surprising common conclusion was that with actions and not signals publicly observable, there was a positive chance that a 'herd' would eventually arise: Everyone after some period would make the identical less profitable decision.

This is a compelling 'pathological' result: Individuals do not eventually learning the true state of the world despite the more than sufficient wealth of information. So let's seriously analyze the model whence it arises. For we believe that learning from others' actions is economically important, and perhaps encompasses the greater part of individual information acquisition that occurs in society at large. As such, from a theoretical standpoint, it merits the scrutiny long afforded the single-person learning results, and the rational expectations literature.

In this paper, we attempt a systematic analysis of the above observational learning paradigm on two fronts: First, we develop a simple dynamic mathematical framework rich enough to describe a rich array of observational learning dynamics. This offers key insights into the probabilistic foundations of observational learning, and that allows us to relatively painlessly generalize the economics at play. Second, we refer to the results of Banerjee (1992) and BHW as the standard herding story, and proceed to spin alternative more economic stories which question the robustness of their pointed conclusion. We in fact find that herding is not the only possible 'pathological' outcome. For not only is it possible that all individuals may almost surely end up taking the correct action, but under just as plausible conditions, social dynamics may well converge to a situation where no one can draw any inference at all from history! We relate this 'confounded learning' outcome to a result due to McLennan (1984) from the single-person experimentation literature. We then argue that the twin pathologies of herding and confounded learning are essentially the only possible ways in which individuals eventually fail to learn the truth.

On Herding and 'Cascades'

While 'herding' certainly has a negative connotation, BHW in fact pointed out that everyone would almost surely eventually settle on a common decision. For this reason, they introduced the arguably more colorful terminology of a *cascade*, referring to any such infinite train of individuals who decide to act irrespective of the content of their

[1] Economic examples include whether to invest in a newly opening and soon to be closing market. A historical instance might have been the decision to enjoy the maiden voyage of the Titanic.

signal. While we find this terminology useful and shall adopt it. we still succomb to the herd in denoting these works and those that followed them the 'herding' literature. The common thread linking these papers is the herding 'pathology' that arises under sequential decision-making when actions and not information signals of previous decision makers are observable. [2]

This definition indirectly rules out the model of Lee (1993), who allows for continuous action spaces which together with a continuous payoff function can perfectly reveal (at least a range of) private signals![3] That no herding arises in Lee's context should come as little surprise — at least to those that read this paper. Indeed, herding pathologies were absent from the rational expectations literature for this very reason,[4] since the Walrasian price can adjust continuously.[5] As a result, with one-dimensional information at least. marginal changes in individuals' private signals all have impact on the publicly-observed price. It is necessary in some sense that the entry of new information be endogenously 'lumpy' for herding to occur (so that it can eventually be choked off altogether).

A Tour of the Paper

We first focus on the key role played by the informational assumptions underlying the standard story. A crucial element is that the individuals cannot get arbitrarily strong private signals. so that their private likelihood ratio is bounded away from zero and infinity. For in that case. a finite history can provide such a strong signal, that even the most doctrinaire individual dare not quarrel with its conclusion. When this "bounded beliefs" assumption is relaxed, incorrect herds cannot arise — and in fact, eventually all individuals will make the correct choice.[6] That result is an application of the Borel-Cantelli Lemma. since if individuals were herding on a wrong action, then there would with probability one appear an individual with so strong a private belief that he would take another action. thereby revealing his very strong information and overturning the herd. Indeed, casual empiricism suggests that individuals who are arbitrarily tenacious in their beliefs do exist. But this assumption is largely a modelling decision, and therein lies its ultimate justification. For it provides us with a richer model than possible in the standard story, allowing us to consider natural economic modifications that hinder informational transmission. and ask if convergence still (almost surely) occurs. This natural approach to robustness was simply not possible in the framework of Banerjee (1992) and BHW.

That a single individual can 'overturn the herd' turns out to be the key insight into the nonexistence of herding with unbounded beliefs. So our first key economic innovation is to prevent this from happening, and introduce noise into the model. Counterintuitively,

[2] It is noteworthy that by this definition, herding was discovered in some veiled form in a concluding example in Jovanovic (1987).

[3] Contrast, for a moment Banerjee's (1992) model which had a continuous action space but a discontinuous payoff function.

[4] To be perfectly clear, we are referring to the literature on dynamic price formation under incomplete information. For a good take on this field. see Bray and Kreps (1987).

[5] This assumes that individuals can continuously adjust their trading quantities. If not, Avery and Zemsky (1992) have shown that a temporary herd may arise.

[6] This idea, not as cleanly expressed, was introduced in Smith (1991), which this paper and Smith and Sorensen (1994) now supersede.

even with a constant inflow of crazy individuals, we still find that learning is complete in the sense that everyone (sane) eventually learns the true state of the world. We then turn to a parallel reason why the actions of isolated individuals need not matter, namely multiple individuals' types. That is, we relax the assumption that all individuals have the same *preferences*. This obviously was a crucial underlying tenet for the herding results of Banerjee (1992) and BHW. For instance, suppose on a highway under construction, depending on how the detours are arranged, that those going to Chicago should merge right (resp. left), with the opposite for those headed toward St. Louis. If one knows that roughly 75% are headed toward Chicago, then absent any strong signal to the contrary, those headed toward St. Louis should take the lane 'less traveled by'. That the resulting dynamics might converge to a totally uninformative inference even with arbitrarily strong private signals is the surprising content of Theorem 8.

We conclude with a brief discussion of costly information and payoff externalities. We do not study endogenous timing, as we have little to add to the recent findings of Chamley and Gale (1992). In a separate more involved work in progress, we investigate what happens when individuals do not perfectly observe the order of previous individuals' moves. This is yet another (more typical) reason for why contrary actions of isolated individuals might have very little effect. Unfortunately, standard martingale results cannot be applied, and therefore it falls outside the scope of this paper.

Identifying the appropriate stochastic processes that are martingales turns out to have been a crucial step in our analysis. The essential analytics of the paper build on the fact that public likelihood ratio is (conditional on the state) a martingale *and* a homogeneous Markov chain. The Markovian aspect of the dynamics allows us (just as it did Futia (1982)) to drastically narrow the range of possible long run outcomes, as we need only focus on the ergodic set. This set is wholly unrelated to initial conditions, and depends only on the transition dynamics of the model. By contrast, the Martingale property of the model — which is unavailable in Futia (1982) — affords us a different glimpse into the long run dynamics, tying them down to the initial conditions in expectation. As it turns out, this allows us to eliminate from consideration the not implausible elements of the ergodic set where everyone entertains entirely false beliefs in the long run.

Section 2 outlines the basic mathematical framework within which we are operating. Section 3 takes a brief mathematical detour, developing some key generic insights on the underlying probabilistic dynamics. We return to the characterization of when herding occurs in section 4, and explore the robustness in sections 5, 6, 7. An appendix, among other things, derives some new results on the local stability of stochastic difference equations. This result, whose absence from the literature (to our knowledge!) greatly surprised us, ought to prove widely applicable across economics and the mathematical (social) sciences in general.

2. THE STANDARD MODEL

2.1 Some Notation

We first introduce a background probability space $(\Omega, \mathcal{E}, \nu)$. This space underlies all random processes in the model, and is assumed to be common knowledge.

An infinite sequence of individuals $n = 1, 2, \ldots$ sequentially takes actions in that exogenous order. Individuals observe the actions of all predecessors. There are two *states of the world* (or more simply, *states*), labelled H ('high') and L ('low'). Formally, this means that the background state space Ω is partitioned into two events Ω^H and Ω^L, called H and L.[7] The results derived below will go through with any finite number of states, but the notation becomes considerably harder. Therefore we stick to the two states case, and later explain carefully how we can modify the analysis to more states. Let the common prior belief be that $\nu(H) = \nu(L) = 1/2$. That individuals have common priors is a standard modelling assumption, see e.g. Harsanyi (1967–68). Also, a flat prior over states is truly WLOG, for it will turn out that more general priors will be formally equivalent to a renormalization of the payoffs, as seen in section 2.2 below.

Individual n receives a private random signal, $\sigma_n \in \Sigma$, about the state of the world. *Conditional on the state,* the signals are assumed to be i.i.d. across individuals. It is common knowledge that in state H (resp. state L), the signal is distributed according to the probability measure μ^H (resp. μ^L). Formally, we mean that $\sigma_n : \Omega \to \Sigma$ is a stochastic variable, and $\mu^H = \nu^H \circ \sigma_n^{-1}$ and $\mu^L = \nu^L \circ \sigma_n^{-1}$, where ν^H (resp. ν^L) is the measure ν conditioned on the event Ω^H (resp. Ω^L). To ensure that no signal will perfectly reveal the state of the world, we shall insist that μ^H and μ^L be mutually absolutely continuous.[8] Consequently, there exists a positive and finite Radon-Nikodym derivative $g = d\mu^H/d\mu^L : \Sigma \to (0, \infty)$ of μ^H w.r.t. μ^L. And to avoid trivialities, we shall *rule out* $g = 1$ almost surely,[9] so that μ^H and μ^L are not the same measure; this will ensure that some signals are informative about the state of the world.

Using Bayes' rule, the individual arrives at what we shall refer to as his *private belief* $p(\sigma) = g(\sigma)/(g(\sigma)+1) \in (0, 1)$ that the state is H. Conditional on the state, private beliefs are i.i.d. across individuals because signals are. In state H (resp. state L), p is distributed with a c.d.f. F^H (resp. c.d.f. F^L) on $(0, 1)$. The distributions F^H and F^L are subtly linked. In Appendix A, we prove among other things that F^H and F^L have the same support,[10] and that $F^L - F^H$ increases (weakly) on $[0, 1/2]$ and decreases (weakly) on $[1/2, 1]$. Denote the common support of F^H and F^L by $\mathrm{supp}(F)$. The structure of $\mathrm{supp}(F)$ will play a major role in the definition of herds. Observe that the common support of F^H and F^L, which we shall denote $\mathrm{supp}(F)$, coincides with the range of $p(\cdot)$ on Σ.[11] It is therefore

[7] For later reference, refer to the restricted sigma fields as $\mathcal{E}^H$ and $\mathcal{E}^L$, respectively.

[8] Recall that μ^H is absolutely continuous w.r.t. μ^L if $\mu^L(S) = 0 \Rightarrow \mu^H(S) = 0 \; \forall S \in \mathcal{S}$, where $\mathcal{S}$ is the σ-algebra on Σ. By the Radon-Nikodym Theorem, there exists then a unique $g \in L^1(\mu^L)$ such that $\mu^H(S) = \int_S g \, d\mu^L$ for every $S \in \mathcal{S}$. See Rudin (1987).

[9] Note that with μ^H and μ^L mutually a.c., "almost sure" assertions are well-defined without specifying which measure.

[10] Recall that the support of a probability measure is any measurable set accorded probability 1. But throughout the paper, 'the' support is well-defined modulo measure zero equivalence.

[11] While the Radon-Nikodym derivative g is only determined with probability one, we can select a version

important to observe that the underlying results are ultimately driven by the probability measures μ^H and μ^L, which are the primitive of the model.

By construction, $\text{co}(\text{supp}(F)) \equiv [\underline{b}, \bar{b}] \subseteq [0, 1]$ with $0 \le \underline{b} < \bar{b} \le 1$.[12] We shall say that the private beliefs are *bounded* if $0 < \underline{b} < \bar{b} < 1$. If $\text{co}(\text{supp}(F)) = [0, 1]$, we simply call the private beliefs *unbounded*.

Each individual can choose from a finite set of actions $\langle a_m, m \in \mathcal{M}\rangle$, where $\mathcal{M} = \{1, \ldots, M\}$. Action a_m has a (common) payoff $u^H(a_m)$ in state H and $u^L(a_m)$ in state L. The objective of the individual is to take the action that maximizes his expected payoff. We assume WLOG that no action is weakly dominated (by all other actions), and to avoid trivialities we insist that at least two such undominated actions exist. Before deciding upon an action, the individual can observe the entire action history. We shall loosely denote the action profile of *any finite number of* individuals as h. Exactly how the individual uses that history is considered in the next subsection.

2.2 Preliminary Results

Action Choice

Given a posterior belief $r \in (0, 1)$ that the state is H, the expected payoff of action a is $ru^H(a) + (1 - r)u^L(a)$. Figure 1 portrays the content of the next result.

Lemma 1 *The interval* $(0, 1)$ *partitions into relatively closed subintervals* $I_1, \ldots, I_M$ *overlapping at endpoints only, such that action* a_m *is optimal when the posterior* $r \in I_m$.

Proof: As noted, the payoff of each action is a linear function of r. Hence, because by the assumption that action a_m is strictly best for some r, there must be a single open subinterval of $(0, 1)$ where action a_m strictly dominates all other actions. That this is a partition follows from the fact that there exists at least one optimal action for each posterior $r \in (0, 1)$. ◇

We now WLOG strictly order the actions such that a_m is optimal exactly when the posterior $r \in [\bar{r}_{m-1}, \bar{r}_m] \equiv I_m$, where $0 = \bar{r}_0 < \bar{r}_1 < \cdots < \bar{r}_M = 1$. Let us further introduce the tie-breaking rule that individuals take action a_m, and not a_{m+1}, whenever $r = \bar{r}_m$. Note that action a_M (resp. a_1) is optimal when one is certain that the state is H (resp. L). Indeed, perfect information leads one to take the correct action, and with more actions than states, we might think of it as one of the 'extreme' actions; however, as decisions are generally taken without the luxury of such focused beliefs, an 'insurance' action may well be chosen.

We can now see how unfair priors are equivalent to a simple payoff renormalization, as asserted earlier. For the characterization in Lemma 1 is still valid, since reference is only made to the posterior beliefs; moreover, the key defining indifference relation $\bar{r}_m u^H(a_m) + (1 - \bar{r}_m)u^L(a_m) = \bar{r}_m u^H(a_{m+1}) + (1 - \bar{r}_m)u^L(a_{m+1})$ implies that

$$\text{posterior odds} = \frac{1 - \bar{r}_m}{\bar{r}_m} = \frac{u^H(a_m) - u^H(a_{m+1})}{u^L(a_m) - u^L(a_{m+1})}$$

of it such that the above holds.

[12] Here, $\text{co}(A)$ denotes the convex hull of the set A.

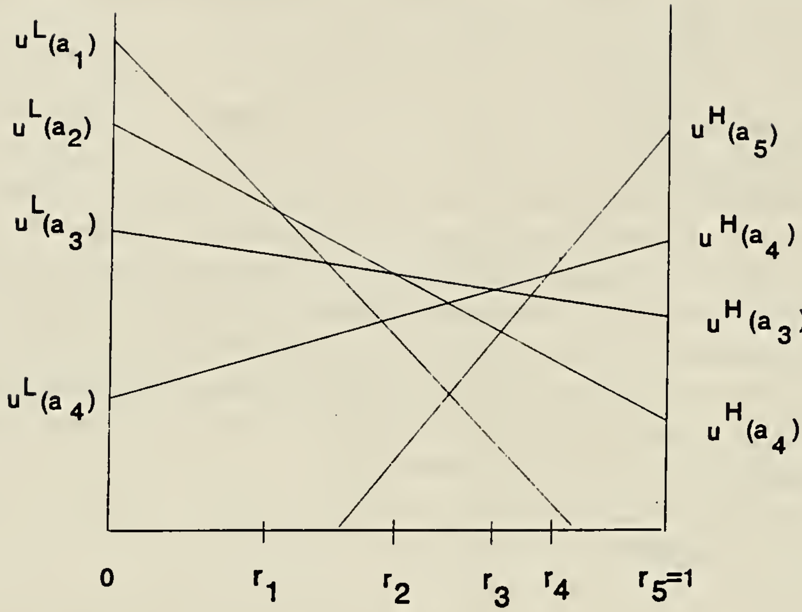

Figure 1: **Example Payoff Frontier.** The diagram depicts the payoff of each of five actions as a function of the posterior belief r that the state is H. The individual simply chooses the action yielding the highest payoff.

Since unfair priors merely serve to multiply the posterior odds by a common constant, the thresholds $\bar{r}_0, \ldots, \bar{r}_M$ are all unchanged if we merely multiply all payoffs in state H by the same constant.

Individual Learning

We now consider how an individual's optimal decision rule incorporates the observed action history and his own private belief. In so doing, we could proceed inductively, and first derive the first individual's decision rule as a function of his private belief; next, we could describe how the second individual bases his decision on the private belief *and* on the first individual's action, and so on. Instead, we shall collapse this reasoning processes, and simply say that individual n knows the decision rules of all the previous agents, and acts accordingly. He can use the common prior to calculate the ex ante (that is, as of time-0) probability of any action profile h in either of the two states. We shall denote these probabilities by $\pi^H(h)$ and $\pi^L(h)$, and let the resulting *likelihood ratio* that the state is L (that is, *low* and not high) be $\ell(h) = \pi^L(h)/\pi^H(h)$. Similarly, let $q(h)$ be the *public belief* that the state is H, i.e.

$$q(h) = \frac{\pi^H(h)}{\pi^H(h) + \pi^L(h)} = 1/(1 + \ell(h)).$$

Think of $q(h)$ as the belief an individual facing the history h would entertain if he had a purely neutral private belief. Given the one-to-one relationship between q and $\ell = (1-q)/q$, we may also loosely refer to the likelihood ratio as the public belief.

A final application of Bayes rule yields the posterior belief r (that the state is H) in terms of the public signal — or equivalently the likelihood ratio $\ell(h)$ — and the private belief p:

$$r = \frac{p\,\pi^H(h)}{p\,\pi^H(h) + (1-p)\,\pi^L(h)} = \frac{p}{p + (1-p)\,\ell(h)} = \frac{1}{1 + \frac{1-p}{p}\ell(h)} \tag{1}$$

Lemma 2 (Private Belief Thresholds) *After history h is observed, there exist threshold values $\bar{p}_m(h) \in (0,1)$, such that a_m is chosen exactly when the private belief satisfies $p \in (\bar{p}_{m-1}(h), \bar{p}_m(h)]$, where $\bar{p}_M(h) = 1$ and for all $m = 0, \ldots, M-1$,*

$$\frac{\bar{p}_m(h)}{1-\bar{p}_m(h)} = \frac{\bar{r}_m}{1-\bar{r}_m}\ell(h) \tag{2}$$

The proof is simple. The thresholds simply come from a well-known reformulation of (1) as posterior odds $(1-r)/r$ equal the private odds $(1-p)/p$ times the likelihood ratio $\ell(h)$. The strict inequalities are consequences of the tie-breaking rule, and the fact that (1) is strictly increasing in p.

Observe that corresponding to $\bar{r}_0(h) = 0$ and $\bar{r}_M(h) = 1$, we have $\bar{p}_0(h) = 0$ and $\bar{p}_M(h) = 1$ after any history h. Later, when referring to (2) and elsewhere, we shall suppress the explicit dependence of $\ell(h)$ on h whenever convenient, and write $\bar{p}_m(\ell)$ instead of $\bar{p}_m(h)$. This is not entirely unjustified because the likelihood ratio is a sufficient statistic for the history. Written as such, $\ell \mapsto \bar{p}_m(\ell)$ is an increasing function.

Corporate Learning

We shall denote the likelihood ratio and public belief confronting individual n as ℓ_n and q_n, respectively.[13] Since the first agent has not observed any history, we shall normalize $\ell_1 = 1$. As signals, and thereby actions, are random, the likelihood ratio $\langle \ell_n \rangle_{n=1}^{\infty}$ and public beliefs $\langle q_n \rangle_{n=1}^{\infty}$ are both stochastic processes, and it is important to understand their long-run behavior.

First, as is standard in learning models,[14] the public beliefs constitute a martingale.[15]

Lemma 3 (The Unconditional Martingale) *The public belief $\langle q_n \rangle$ is a martingale, unconditional on the state of the world.*

Proof: Individual n's action only depends on the history through ℓ_n, or equivalently $q_n = 1/(1+\ell_n)$. Think of his private belief as being realized after this observation. Ex ante to this realization, let $\alpha_m^H(q_n)$ (resp. $\alpha_m^L(q_n)$) be the conditional probability that action a_m is taken in state H (resp. state L). Then the conditional expectation of the next public belief is

$$E[q_{n+1} \mid q_1, \ldots, q_n] = E[q_{n+1} \mid q_n]$$

[13]Throughout the paper, m subscripts will denote actions, and n subscripts individuals.

[14]For instance, Aghion, Bolton, Harris and Jullien (1991) establish this result for the experimentation literature.

[15]We really ought to specify the accompanying sequence of σ-algebras is the stochastic process, in order to speak about a martingale; however, these will be suppressed because they are simply the ones generated by the process itself.

$$= q_n \sum_{m\in\mathcal{M}} \alpha_m^H(q_n) \frac{1}{1+\ell_n \frac{\alpha_m^L(q_n)}{\alpha_m^H(q_n)}} + (1-q_n) \sum_{m\in\mathcal{M}} \alpha_m^L(q_n) \frac{1}{1+\ell_n \frac{\alpha_m^L(q_n)}{\alpha_m^H(q_n)}}$$

$$= q_n \sum_{m\in\mathcal{M}} \alpha_m^H(q_n) \frac{q_n\alpha_m^H(q_n) + (1-q_n)\alpha_m^L(q_n)}{q_n\alpha_m^H(q_n) + (1-q_n)\alpha_m^L(q_n)}$$

$$= q_n$$

◇

This martingale describes the forecast of subsequent public beliefs by individuals in the model, since they do not, of course, know the true state of the world: Prior to receiving his signal, individual n's best guess of the public belief that will confront his successor is the current one. But for our purposes, an unconditional martingale tells us little about convergence. For that, we must condition on the state of the world, and it is well-known that will render the public belief $\langle q_n \rangle$ a *sub*martingale in state H (and a *super*martingale in state L), i.e. $E[q_{n+1} \mid H, q_1, \ldots, q_n] \geq q_n$. This will follow from Lemma 4 below. Essentially, the public beliefs are expected to become weakly more focused on the true state of the world — a result much weaker than we seek.

Given that $\langle q_n \rangle$ is expected to increase in state H, we at least find it rather surprising (if easy to prove) that $\langle \ell_n \rangle = \langle (1+q_n)/q_n \rangle$ remains constant in expectation.

Lemma 4 (The Conditional Martingale) *Conditional on the state of the world H (resp. state L), the likelihood ratio $\langle \ell_n \rangle$ (resp. $\langle 1/\ell_n \rangle$) is a martingale.*

Proof: Given the value of ℓ_n, the next individual will take one of the available actions, depending on his prior belief. In state H, action a_m is taken with some conditional probability $\alpha_m^H(\ell_n)$, while in state L, the chance is $\alpha_m^L(\ell_n)$. Thus, the conditional expectation of the likelihood ratio in state H is

$$E[\ell_{n+1} \mid H, \ell_1, \ldots \ell_n] = E[\ell_{n+1} \mid H, \ell_n] = \sum_{m\in\mathcal{M}} \alpha_m^H(\ell_n)\, \ell_n \, \frac{\alpha_m^L(\ell_n)}{\alpha_m^H(\ell_n)} = \ell_n \sum_{m\in\mathcal{M}} \alpha_m^L(\ell_n) = \ell_n$$

◇

Lemma 5 *In state H (resp. state L), the likelihood ratio $\langle \ell_n \rangle$ (resp. $\langle 1/\ell_n \rangle$) converges almost surely to a limiting stochastic variable which takes on values in $[0, \infty)$.*

Proof: This follows directly from the Martingale Convergence Theorem, as found for instance in Breiman (1968), Theorem 5.14. ◇

That the limiting likelihood ratio (if it exists) cannot place positive weight on ∞ was clear anyway clear, for the martingale property yields $E[\ell_\infty \mid H, \ell_1] = \ell_1 = 1$. This crucially precludes individuals (eventually) being wholly mistaken about the true state of the world.

In the sequel, our goals are two-fold. Suppose the state is H. First, we wish to establish general conditions guaranteeing that $\ell_n \to 0$, so that all individuals with unconcentrated beliefs eventually learn the true state of the world. Whenever ℓ is very close to 0, only individuals with very strong signals will take a suboptimal action. Second, we also wish to prove that eventually all individuals will almost surely take the optimal action. While such a result is perhaps more straightforward, it does not lend itself to the more general framework we shall later consider.

More States and Actions

The analysis goes through virtually unchanged with a denumerable action space. Rather than a finite partition of $[0, 1]$ in Lemma 1. we get a countable partition. and thus a countable set of posterior belief thresholds $\bar{r}$.[16] In this way, Lemma 2 will yield the threshold functions $\bar{p}$ just as above. The martingale properties of the model are preserved.

We can also handle any finite number S of states. Given pairwise mutually absolutely continuous measures μ^s for each state, we could fix one reference state, and use it to define $S-1$ likelihood ratios. Again, each likelihood ratio would be a convergent conditional martingale. The complication is largely notational, as the optimal decision rules become rather cumbersome. Rather than the simple partitioning of $[0, 1]$ into closed subintervals. we would now have a unit simplex in $\mathbb{R}^{S-1}$ sliced into closed convex polytopes. We leave it to the reader to ponder the optimal notation, but we simply assert that the above results would still obtain.

Herding, Cascades, and Complete Learning

We are now positioned to define some fundamental concepts. We find it best to think of all dynamics as occurring on two different levels. From an observational point of view, we wish to use the popular street language of a *herd*, as adopted in Banerjee (1992). Say that a *herd arises* if for some n all agents starting at the nth choose the same action.

But in the more general framework with multiple types and noise that we soon consider, herds need not arise and yet convergence in beliefs may still obtain. For this notion we first appeal to BHW's term *cascade.* We say that a *cascade arises* if for after some stage n, $\text{supp}(F) \subseteq (\bar{p}_{m-1}(\ell_n), \bar{p}_m(\ell_n)]$ for some m. But even this notion is not sufficient for our purposes. Adopting the terminology introduced in Aghion et al. (1991), we shall call the learning *complete* so long as individuals' posterior beliefs eventually become arbitrarily focused on the true state of the world: that is, if the interval $(\bar{p}_{m-1}(\ell_n), \bar{p}_m(\ell_n)]$ converges to a set that contains $\text{supp}(F)$ as $n \to \infty$, where action a_m is optimal. Otherwise. if posterior beliefs do not eventually become arbitrarily focused on the true state of the world. then we shall say that learning is *incomplete.* Observe that complete learning will not imply a cascade with unbounded beliefs, for they exist some individuals with signals so strong as to not wish to ignore them; conversely, if there is a cascade on the optimal action, then complete learning obtains.

It is easy to see the equivalence of cascades and herds. Indeed, if a cascade on action a_m arises at stage n in the above sense, then by Lemma 2, individual $n+1$ will (irrespective of the state) necessarily take action a_m; therefore, $\ell_{n+1} = \ell_n$, and so a cascade on action a_m exists at stage $n+1$. Thus, all private belief thresholds are unchanged by (2), and $\text{supp}(F) \subseteq (\bar{p}_{m-1}(\ell_{n+1}), \bar{p}_m(\ell_{n+1})]$ too. The original intuition of BHW or Banerjee (1992) obtains: Each individual takes an action which does not reveal any of his private information, and so the public belief is unchanged.

[16]This may mean that we cannot necessarily well order the order the belief thresholds, nor as a result the actions.

3. DISCRETE DYNAMICAL SYSTEMS

Before tackling the main theorems, we shall step back from the model, and consider a mathematical abstraction that will encompass the later variations. The general framework that we introduce includes, but is not confined to, the evolution of the likelihood ratio $\langle \ell_n \rangle$ over time viewed as a stochastic difference equation.[17]

The context is as follows. Given is a finite set $\mathcal{M}$, and functions $\varphi(\cdot\,,\,\cdot) : \mathcal{M} \times \mathbb{R}_+ \to \mathbb{R}_+$, and $\psi(\cdot \mid \cdot) : \mathcal{M} \times \mathbb{R}_+ \to [0,1]$ meeting two restrictions. First, $\psi(\cdot \mid \ell)$ must be a probability measure for all $\ell \in \mathbb{R}_+$, or

$$\sum_{m \in \mathcal{M}} \psi(m|\ell) = 1.$$

Second, the following 'martingale property' must hold for all $\ell \in \mathbb{R}_+$:

$$\sum_{m \in \mathcal{M}} \psi(m|\ell)\ \varphi(m,\ell) = \ell \tag{3}$$

Finally, equip $\mathbb{R}_+ = [0,\infty)$ with the Borel σ-algebra $\mathcal{B}$, and define a transition probability $P : \mathbb{R}_+ \times \mathcal{B} \to [0,1]$ as follows:

$$P(\ell, B) = \sum_{m \mid \varphi(m,\ell) \in B} \psi(m|\ell) \tag{4}$$

for any $B \in \mathcal{B}$. For our immediate application, one can think of $\psi(m|\ell)$ as the chance that the next agent takes action m when faced with likelihood ℓ, and $\varphi(m,\ell)$ as the resulting continuation likelihood ratio.

Suppose for definiteness that we are given a (measurable) Markov stochastic process $\langle \ell_n \rangle_{n=1}^{\infty}$ on $(\Omega^H, \mathcal{E}^H, \nu^H)$, where for each n, $\ell_n : \Omega^H \to \mathbb{R}_+$. Transition from ℓ_n to ℓ_{n+1} is described by the transition probability P. We assume that $E\ell_1 < \infty$; in applications we shall always assume that ℓ_1 is identically 1, so this is not restrictive.[18] Denote by $\mathcal{F}_n$ the σ-field in $(\Omega^H, \mathcal{E}^H)$ generated by $(\ell_1, \ldots, \ell_n)$. Clearly, ℓ_n is $\mathcal{F}_n$-measurable, and it follows from (3) that $\langle \ell_n, \mathcal{F}_n \rangle$ is actually a martingale,[19] thus justifying our earlier casual description of property (3). Indeed,

$$E[\ell_{n+1}|\ell_1, \ldots, \ell_n] = E[\ell_{n+1}|\ell_n] = \int_{\mathbb{R}_+} tP(\ell_n, dt) = \sum_{m \in \mathcal{M}} \psi(m|\ell_n)\varphi(m,\ell_n) = \ell_n$$

Since $\langle \ell_n \rangle$ is a martingale on $\mathbb{R}_+$, we know from the Martingale Convergence Theorem that it converges almost surely in $\mathbb{R}_+$. Denote the limiting stochastic variable by $\tilde{\ell}$. We now characterize the limit.

[17] Arthur, Ermoliev and Kaniovski (1986) consider a stochastic system with a seemingly similar structure — namely, a 'generalized urn scheme'. Their approach, however, differs fundamentally from ours insofar as here it is of importance not only how many times a given action has occurred, but exactly *when* it occurred. But while we cannot apply their results, we owe them a debt of inspiration.

[18] Notice that the system has a discrete transition function; therefore, if ℓ_1 has a discrete distribution the process will be a discrete (in fact, countably infinite) Markov chain. One might think that it would be possible to apply standard results about the convergence of discrete Markov chains, but in fact such results are not useful here. While the state space is certainly countable, all states (which will soon be interpreted as likelihood functions) are in general transitory, and so standard results are useless.

[19] No ambiguity arises if we simply say that $\langle \ell_n \rangle$ is a martingale.

Theorem 1 (Stationarity) *Assume that for all $m \in \mathcal{M}$, the two functions $\ell \mapsto \varphi(m,\ell)$ and $\ell \mapsto \psi(m|\ell)$ are continuous. Suppose that $\ell_n \to \tilde{\ell}$ almost surely. Then for all $m \in \mathcal{M}$ and for all $\ell \in \operatorname{supp}(\tilde{\ell})$, stationarity obtains, i.e.*

$$\psi(m|\ell) > 0 \ \Rightarrow \ \varphi(m,\ell) = \ell \tag{5}$$

As this theorem will follow from the next theorem, its proof is deferred. That implication (5) is truly a stationarity condition is best seen — by means of (4) — in its alternative formulation $P(\ell, \{\ell\}) = 1$.

The intuition behind Theorem 1 is rather simple. Since ℓ_n converges almost surely to $\tilde{\ell}$, it also converges weakly (in distribution) to $\tilde{\ell}$. As the process is also a Markov chain, it is intuitive that the limiting distribution is invariant for the transition P, as described in Futia (1982). In fact, we can prove Theorem 1 along these lines, but the continuity assumptions are subtly hard-wired into the final stage of the argument to prove that the limiting distribution is invariant. As we wish to do away with continuity, we establish an even stronger result. Motivated by the fact that (5) is violated for m exactly when neither $\psi(m|\ell)$ nor $\varphi(m,\ell) - \ell$ is zero, we have

Theorem 2 (Generalized Stationarity) *Assume that the open interval $I \subseteq \mathbb{R}_+$ has the property*

$$\exists \varepsilon > 0 \ \forall \ell \in I \ \exists m \in \mathcal{M} : \ \psi(m|\ell) > \varepsilon, \ |\varphi(m,\ell) - \ell| > \varepsilon \tag{$\star$}$$

Then I cannot contain any point from the support of the limit, $\tilde{\ell}$.

Proof: Let I be an arbitrary open interval satisfying ($\star$) for $\varepsilon > 0$, and suppose by way of contradiction that there exists $\bar{\ell} \in I \cap \operatorname{supp}(\tilde{\ell})$. Let $J = (\bar{\ell} - \varepsilon/2, \bar{\ell} + \varepsilon/2) \cap I$. By ($\star$), for all $\ell \in J$, there exists $m \in \mathcal{M}$ such that $\psi(m|\ell) > \varepsilon$, and $\varphi(m,\ell) \notin J$. Because $\bar{\ell} \in \operatorname{supp}(\tilde{\ell})$, there is positive probability that $\ell_n \in J$ eventually. But whenever $\ell_n \in J$, there is a probability of at least ε that $\ell_{n+1} \notin J$. Since $\langle \ell_n \rangle$ is a Markov process, the events $\{\ell_n \in J, \ell_{n+1} \notin J\}_{n=1}^{\infty}$ are independent. Thus, by the (second) Borel-Cantelli Lemma, the process $\langle \ell_n \rangle$ must almost surely leave J infinitely often — contradicting the claim that with positive probability it is eventually in J. Hence, $\bar{\ell}$ cannot exist. $\diamondsuit$

Corollary *Assume that $\bar{\ell} \in \operatorname{supp}(\tilde{\ell})$. Then for each $m \in \mathcal{M}$, either $\ell \mapsto \varphi(m,\ell)$ or $\ell \mapsto \psi(m|\ell)$ is discontinuous at $\bar{\ell}$, or the stationarity condition* (5) *obtains.*

Proof: If there is an m such that $\bar{\ell}$ does not satisfy (5) and both $\ell \mapsto \varphi(m,\ell)$ and $\ell \mapsto \psi(m|\ell)$ are continuous, then there is an open interval I around $\bar{\ell}$ in which $\psi(m|\ell)$ and $\varphi(m,\ell) - \ell$ are both bounded away from 0. This implies that ($\star$) obtains, and so Theorem 2 yields an immediate contradiction. $\diamondsuit$

Finally, it is obvious that Corollary implies Theorem 1.

More States and Actions

Once again, we could easily have handled a countable action space $\mathcal{M}$, as the finiteness of $\mathcal{M}$ was never used. Also, given any finite number $S > 2$ states of the world, all results still obtain. For $\langle \ell_n \rangle$ would then be a stochastic process in $\mathbb{R}^{S-1}$, and we need only refer to the open intervals I and J in Theorem 2 (and its proof) as open balls.

4. THE MAIN RESULTS

We are now ready to characterize exactly when either cascades or herding arises. To do so, we shall recast the model of section 2 in the language of section 3. Fix the likelihood ratio ℓ, and assume WLOG that the state is H. By Lemma 2, the individual takes action a_m exactly when his private belief is in the interval $(\bar{p}_{m-1}(\ell), \bar{p}_m(\ell)]$. Since this occurs with chance $F^H(\bar{p}_m(\ell)) - F^H(\bar{p}_{m-1}(\ell))$ in state H, and with chance $F^L(\bar{p}_m(\ell)) - F^L(\bar{p}_{m-1}(\ell))$ in state L, we have

$$\psi(m|\ell) = F^H(\bar{p}_m(\ell)) - F^H(\bar{p}_{m-1}(\ell)) \tag{6}$$

$$\varphi(m,\ell) = \ell\frac{F^L(\bar{p}_m(\ell)) - F^L(\bar{p}_{m-1}(\ell))}{F^H(\bar{p}_m(\ell)) - F^H(\bar{p}_{m-1}(\ell))} \tag{7}$$

in the notation of section 3.

We know from Lemma 5 that $\ell_\infty = \lim_{n\to\infty} \ell_n$ almost surely exists. We now apply Theorem 2 to get a precise characterization of the limiting stochastic variable. Recall that $\mathrm{co}(\mathrm{supp}(F)) = [\underline{b}, \bar{b}]$. The crucial question is whether the individuals can have arbitrarily informative private signals or not, i.e. whether $[\underline{b}, \bar{b}] = [0, 1]$ or $[\underline{b}, \bar{b}] \subset (0, 1)$.

4.1 Bounded Beliefs

Assume that the private beliefs are bounded. Our approach is two-fold. We first exhibit 'action absorbing basins', each corresponding to an action choice, in which all learning stops, and individuals act irrespective of their signals. We then argue that in fact the dynamics eventually almost surely end up in one of these basins, i.e. that cascades must occur.

Lemma 6 (Action Absorbing Basins) *There are (possibly empty) intervals $J_1, \ldots, J_M$ in $[0, \infty)$, where $J_m = \{\ell \mid [\bar{p}_{m-1}(\ell), \bar{p}_m(\ell)] \subseteq \mathrm{supp}(F)\}$, such that almost surely when $\ell \in J_m$ the individual takes action a_m, and the next likelihood ratio will still be in J_m. Moreover,*
(1) not all intervals are empty, as $J_1 = [\bar{\ell}, \infty)$ and $J_M = [0, \underline{\ell}]$ for some $0 < \underline{\ell} < \bar{\ell} < \infty$;
(2) the intervals have disjoint interiors, and are in fact inversely ordered in the sense that all elements of J_{m_2} are strictly smaller than any element of J_{m_1} when $m_2 > m_1$.

Proof: Since $\bar{p}_m(\ell)$ is increasing in m by Lemma 2, $[\bar{p}_{m-1}(\ell), \bar{p}_m(\ell)]$ is an interval for all ℓ. Then J_m is the closure all ℓ that fulfill

$$\bar{p}_{m-1}(\ell) \leq \underline{b} \quad \text{and} \quad \bar{p}_m(\ell) \geq \bar{b} \tag{8}$$

Then disjointness is obvious. Next, if $J_m \neq \varnothing$ then $F^H(\bar{p}_{m-1}(\ell)) = 0$ and $F^H(\bar{p}_m(\ell)) = 1$ for all $\ell \in J_m$. The individual will choose action a_m a.s., and so no updating occurs; therefore, the continuation value is a.s. ℓ, as required.

With bounded beliefs, it is clear that we can always ensure one of the inequalities in (8) for some ℓ, but simultaneously attaining the two may well be impossible. As Lemma 2 yields $\bar{p}_0(\ell) \equiv 0$ and $\bar{p}_M(\ell) \equiv 1$ for all ℓ, it follows that we must have $J_M = [0, \underline{\ell}]$ and $J_1 = [\bar{\ell}, \infty)$, where $0 < \underline{\ell} < \bar{\ell} < \infty$ satisfy $\bar{p}_{M-1}(\underline{\ell}) = \underline{b}$ and $\bar{p}_1(\bar{\ell}) = \bar{b}$.

Finally, let $m_2 > m_1$, with $\ell_1 \in J_{m_1}$ and $\ell_2 \in J_{m_2}$. Then

$$\bar{p}_{m_2-1}(\ell_1) \geq \bar{p}_{m_1}(\ell_1) \geq \bar{b} > \underline{b} \geq \bar{p}_{m_2}(\ell_2) \geq \bar{p}_{m_2-1}(\ell_2)$$

and so $\ell_2 < \ell_1$ because $\bar{p}_{m_2-1}$ is strictly increasing in ℓ. ◇

By rearranging an expression like (2), one can show that ℓ satisfies (8) precisely when

$$\bar{r}_{m-1} \leq \frac{\underline{b}}{\underline{b} + (1-\underline{b})\ell} \quad \text{and} \quad \frac{\bar{b}}{\bar{b} + (1-\bar{b})\ell} \leq \bar{r}_m$$

This can surely also obtain for nonextreme (insurance) actions, and is less likely the smaller is $\underline{b}$, the larger is $\bar{b}$, and the smaller is the interval $[\bar{r}_{m-1}, \bar{r}_m]$.

Theorem 3 (Cascades) *Assume the private beliefs are bounded, and let $\ell_n \to \bar{\ell}$. Then $\bar{\ell} \in J_1 \cup \cdots \cup J_M$ almost surely.*

Proof: Suppose by way of contradiction that $\ell_n \to \bar{\ell} \notin J_1 \cup \cdots \cup J_M$ with positive probability. Assume WLOG the state is H. Then for some m we have $0 < F^H(\bar{p}_m(\bar{\ell})-) < 1$, so that individuals will strictly prefer to choose action a_m for some private beliefs and a_{m+1} for others. Consequently, $\bar{p}_m(\bar{\ell}) > \underline{b}$, and since $\bar{p}_0(\bar{\ell}) = 0 < \underline{b}$, the least such m satisfying $\bar{p}_m(\bar{\ell}) > \underline{b}$ is well-defined. So we may assume $F^H(\bar{p}_{m-1}(\bar{\ell})-) = 0$.

Next, $F^H(\bar{p}_m(\ell)) > 0$ in a neighborhood of $\bar{\ell}$. There are two possibilities:

Case 1. $F^H(\bar{p}_m(\bar{\ell})) > F^H(\bar{p}_{m-1}(\bar{\ell}))$.

Here, there will be a neighborhood around $\bar{\ell}$ where $F^H(\bar{p}_m(\ell)) - F^H(\bar{p}_{m-1}(\ell)) > \varepsilon$ for some $\varepsilon > 0$. We see from (6) that in this neighborhood $\psi(m|\ell)$ is bounded away from 0, while (7) reduces to $\varphi(m, \ell) = \ell F^L(\bar{p}_m(\ell))/F^H(\bar{p}_m(\ell))$, which is also bounded away from $\bar{\ell}$ for ℓ in a neighborhood of $\bar{\ell}$. Indeed, $\bar{p}_m(\bar{\ell})$ is in the interior of co(supp(F)), and so Lemma A.2 guarantees us that $F^L(\bar{p}_m(\ell))$ is bounded above and away from $F^H(\bar{p}_m(\ell))$ for ℓ near $\bar{\ell}$ (recall that $\bar{p}_m$ is continuous). By Theorem 2, $\bar{\ell} \in \text{supp}(\tilde{\ell})$ therefore cannot occur.

Case 2. $F^H(\bar{p}_m(\bar{\ell})) = F^H(\bar{p}_{m-1}(\bar{\ell}))$.

This can only occur if F^H has an atom at $\bar{p}_{m-1}(\bar{\ell}) = \underline{b}$, and places no weight on $(\underline{b}, \bar{p}_m(\bar{\ell})]$. It follows from $F^H(\bar{p}_{m-1}(\bar{\ell})-) = 0$ and $\bar{p}_{m-2} < \bar{p}_{m-1}$, that $F^H(\bar{p}_{m-2}(\ell)) = 0$ for all ℓ in a neighborhood of $\bar{\ell}$. Therefore, $\psi(m-1|\ell)$ and $\varphi(m-1, \ell) - \ell$ are bounded away from 0 on an interval $[\bar{\ell}, \bar{\ell} + \eta)$, for some $\eta > 0$. On the other hand, the choice of m ensures that $\psi(m|\ell)$ and $\varphi(m, \ell) - \ell$ are bounded away from 0 on an interval $(\bar{\ell} - \eta', \bar{\ell}]$, for some $\eta' > 0$. So, once again Theorem 2 (observe the order of the quantifiers!) proves that $\bar{\ell} \notin \text{supp}(\tilde{\ell})$. ◇

Theorem 4 (Herds) *Assume the private beliefs are bounded. Then a herd on some action will almost surely arise in finite time. Absent extreme belief thresholds $\bar{r}_1$ and $\bar{r}_M$, the herd can arise on an action other than the most profitable one.*

Proof: First note that if $\bar{r}_1 > \bar{b}$ (resp. $\bar{r}_M < \underline{b}$) then the first and thus all subsequent individuals a.s. ignore their private signals and take action a_1 (resp. a_M). Now suppose this does not occur, and assume WLOG the state is H. Whenever $\ell \notin [0, \underline{\ell}]$, we know that $F^H(\bar{p}_{M-1}(\ell)) > 0$ so that some action other than a_M is taken with positive probability.

Claim 1: With positive chance, $\ell_n \in J_M$ in a fixed finite number of steps. Consider the following 'fastest ascent' of the likelihood ratio. Suppose that whenever two or

more actions can be taken with positive probability, private beliefs are such that the lowest numbered action is taken. This will have the effect of pushing the public belief toward state L. Then the likelihood will evolve according to $\ell_{n+1} = \ell_n F^L(\bar{p}_m(\ell_n))/F^H(\bar{p}_m(\ell_n)) > \ell_n$. But this can happen only a finite number of times before $\ell_n > \bar{\ell}$. This follows from Lemma A.3, and the fact that $\ell_n \in [\underline{\ell}, \bar{\ell}]$, and so $\bar{p}_m(\ell_n) \in [\bar{p}_m(\underline{\ell}), \bar{p}_m(\bar{\ell})) \subset (0, 1)$. Indeed, we have

$$\ell_{n+1} = \ell_n F^L(\bar{p}_m(\ell_n))/F^H(\bar{p}_m(\ell_n)) \geq \ell_n(1 - \bar{p}_m(\ell_n))/\bar{p}_m(\ell_n),$$

which proves that the step size is bounded below. So, if the likelihood ratio does not start in $[0, \underline{\ell}]$ then it ends up there with a probability strictly less than 1.

Claim 2: $\ell_n \in J_1 \cup \cdots \cup J_M$ almost surely in finite time.
Because $\langle \ell_n \rangle$ is not simply a finite state Markov chain, Theorem 3 does not immediately imply convergence in finite time — for we could conceivably have $\ell_n \to J_1 \cup \cdots \cup J_M$ but $\ell_n \notin J_1 \cup \cdots \cup J_M$ for all n. But in fact this cannot occur, because if the convergence took an infinite number of steps, then the (second) Borel-Cantelli Lemma would imply that the 'upcrossing' of Claim 1 would happen sooner or later, as the events $\{\ell_{n+1} > \ell_n\}_{n=1}^{\infty}$ are independent, conditional on $\sim (J_1 \cup \cdots \cup J_M)$ and on the state of the world H. $\diamondsuit$

So the bottom line is that all individuals eventually stop paying heed to their private signals, at which point the herd begins. Furthermore, herds are either 'correct' or 'incorrect', and arise precisely because it is common knoledge that there are no private beliefs strong enough to overturn the public belief. This is essentially the major pathological learning result obtained by Banerjee (1992) and BHW, albeit extended to $M > 2$ actions.[20]

While we do not assert how fast the convergence occurs, it is easy to see that for ℓ outside $J_1 \cup \cdots \cup J_M$, $\langle \log(\ell_n) \rangle$ follows a random walk, albeit on a countable state space. Still, with absorbing barriers after a fixed number of the same parity jump, results in Billingsley (1986), pp. 128–130, will imply that convergence must be exponentially fast.

4.2 Unbounded Beliefs

Next we present the counterpart to Theorem 3 that was not considered in Banerjee (1992) and BHW. Strictly bounded beliefs turns out to have been the mainstay for their striking pathological herding results. Then

Theorem 5 (Complete Learning) *If the private beliefs are unbounded then almost surely $\ell_n \to 0$ in state H, and $\ell_n \to \infty$ in state L.*

Proof: As usual, let $\tilde{\ell}$ denote the limit of $\langle \ell_n \rangle$ and assume WLOG the state is H. As Lemma 5 tells us that $\text{supp}(\tilde{\ell}) \in [0, \infty)$, it suffices to prove that

Claim: $\text{supp}(\tilde{\ell}) \cap (1/N, N) = \varnothing$ for any natural number $N > 1$.
Let $I_N = (1/N, N)$. First note that with unbounded private beliefs, $\psi(1|\ell) = F^H(\bar{p}_1(\ell))$ is bounded away from 0 on I_N. Next recall that by Lemma A.2, $F^L(r) - F^H(r)$ is increasing on $[0, 1/2)$ and decreasing on $(1/2, 1]$; therefore, $F^L(\bar{p}_1(\ell)) - F^H(\bar{p}_1(\ell))$ is bounded away from 0 on I_N, as is

$$\varphi(1, \ell) - \ell = \ell \left[\frac{F^L(\bar{p}_1(\ell))}{F^H(\bar{p}_1(\ell))} - 1 \right]$$

[20]The analysis of BHW also handled more states. We soon address this generalization.

It then follows from Theorem 2 that I_N does not contain any point from $\text{supp}(\tilde{\ell})$. Finally, let $N \to \infty$ to prove that $\text{supp}(\tilde{\ell}) = 0$. $\diamond$

So if beliefs are unbounded, then eventually everyone becomes 'convinced' that of the true state of the world. That is, it becomes ever harder for a partially revealing private signal to induce an individual to take any other action than the optimal one. Crucially observe that *belief cascades cannot possibly arise with unbounded beliefs*. This follows both from the earlier definition, and the fact that ignoring arbitrarily focused private signals cannot possibly be an optimal policy. Consequently, we cannot (yet) preclude that an infinite *sub*sequence of individuals may get a string of sufficiently perverse signals to lead each to take a suboptimal action.

It is noteworthy that whenever an individual takes a contrary action, subsequent individuals have no choice but to conclude that his signal was very strong, and this is reflected in a draconian revision of the public belief. We say that the herd has been *overturned* by the unexpected action. We shall more carefully formulate this as the *overturning principle*, as it proves central to an understanding of the observational learning paradigm. Assume individual n chooses action a_m. Then individual $n+1$ should, *before* he gets his own private signal, find it optimal to choose action a_m because he knows no more than individual n, and because it is common knowledge that n rationally chose a_m. So, the likelihood ratio after individual n's action, $\ell(h, a_m)$, satisfies

$$\pi^H(h, a_m) = \frac{1}{1 + \ell(h, a_m)} \in (\bar{r}_{m-1}, \bar{r}_m],$$

which is the content of the next lemma.

Lemma 7 (The Overturning Principle) *For any history h, if an individual optimally takes action a_m, then the updated likelihood ratio must satisfy*

$$\ell(h, a_m) \in \left[\frac{1 - \bar{r}_m}{\bar{r}_m}, \frac{1 - \bar{r}_{m-1}}{\bar{r}_{m-1}}\right) \tag{9}$$

The proof is found in Appendix B.

Together with Theorem 5, the overturning principle implies that herds in fact do occur — but only of the nonpathological variety.

Theorem 6 (Correct Herds) *If the private beliefs are unbounded, then almost surely all individuals eventually take action the optimal action.*

Proof: Assume WLOG that the state is H, so that a_M is optimal. Theorem 5 asserts that $\ell_n \to 0$ a.s., and so ℓ_n is eventually in the neighborhood $[0, \frac{1-\bar{r}_{M-1}}{\bar{r}_{M-1}})$ of 0. But by Lemma 7, whenever any other action than action a_M is taken, we exit that neighborhood. $\diamond$

More States and Actions

The convergence result Theorems 3 and 5 do not depend on the action space being denumerable. In the proof of Theorem 3, a technical complication arises, as our choice of the least m such that $\bar{p}_m(\ell) > \underline{b}$ was well-defined because there were only finitely many actions. Otherwise, we could instead just pick m so that $\bar{p}_m$ is close enough to $\underline{b}$ such that

all the "bounded away" assertions hold. Similarly, in the proof of Theorem 5, we could substitute a minimum action threshold $\bar{p}_1$ by one that is arbitrarily close to 0.

Complications are more insidious when it comes to Theorems 4 and 6. First note that with $M = \infty$, both results still obtain without any qualifications provided a unique action is optimal for posteriors sufficiently close to 0 and 1, for then the overturning principle is still valid near the extreme actions. But otherwise, we must change our tune. For instance, with Theorem 6, there may exist an infinite sequence of distinct optimal 'insurance' action choices made such that the likelihood ratio nonetheless converges. This obviously requires that the optimality intervals I_m shrink to a point, which robs the overturning argument of its strength. Yet this is not a serious nonrobustness critique, because the payoff functions of the actions taken by individuals must then converge!

By contrast, incorporating more than two states of the world is rather simple, and the modifications outlined at the end of section 2 essentially apply here too.

5. NOISE

We now turn to the economic robustness of the existing theory, by striking at its central underpinnings. The key role played by the overturning principle is in many ways unsettling: It does not seem 'reasonable' that such large weight be afforded the observation of a single individual's action. For this reason, we first introduce noise into the system, whereby a small fixed flow of individuals either deliberately (that is, they are a 'crazy' type), or by accident (i.e. they 'tremble') do not choose their optimal action. Consequently, no action will have drastic effects, simply because the 'unexpected' is really expected to happen every now and then.

Two theses then seem plausible at this point:

1. The statistically constant nature of noisy individuals does not jeopardize the learning process of the rational informed individuals in the long run, as it can be filtered out: If the likelihood ratio has a trend towards zero without the noise, that trend will be preserved as the underlying force even with the additional noise.

2. The learning will be incomplete, as the stream of isolated crazy individuals making contrary choices will eventually be indistinguishable from the background noise, and the public belief will thus not tend to an extreme value.

We show in this section that in fact the first intuition is correct, so that Theorems 3 and 5 will still hold. We then turn to the more thorny issue of herding.

Two Forms of Noise

Just to be clear, we assume that whether an individual is noisy is not public information, and is distributed independently across individuals.

Craziness. We first posit the existence of crazy individuals in the model. Assume that with probability τ_m, individual n will always take action a_m, regardless of history. To avoid trivialities, we assume a positive fraction $\tau = 1 - \sum_{m=1}^{M} \tau_m > 0$ of 'sane' individuals. In the language of section 3, the dynamics of the likelihood ratio in state H are now described

as follows:

$$\psi(m|\ell) = \tau_m + \tau[F^H(\bar{p}_m(\ell)) - F^H(\bar{p}_{m-1}(\ell))] \tag{10}$$

$$\varphi(m, \ell) = \ell\frac{\tau_m + \tau[F^L(\bar{p}_m(\ell)) - F^L(\bar{p}_{m-1}(\ell))]}{\tau_m + \tau[F^H(\bar{p}_m(\ell)) - F^H(\bar{p}_{m-1}(\ell))]} \tag{11}$$

Trembling. In the second manifestation of noise, all individuals are rational, but some may 'tremble', in the sense of Selten (1975). In particular, individuals randomly take a suboptimal action with probability $\tau(\ell)$ when the likelihood ratio is ℓ; for simplicity, assume that in this event, all other $M-1$ actions are equally likely. With $M = 2$, individuals' actions are wholly uninformative when $\tau(\ell) = 1/2$, so assume that $\tau(\ell)$ is boundedly smaller than 1/2. On the other hand, to avoid completely trivializing the noise, we further insist that $\tau(\ell)$ be bounded away from 0. In state H the dynamics are now

$$\begin{aligned}\psi(m|\ell) &= [1-\tau(\ell)]\left[F^H(\bar{p}_m(\ell)) - F^H(\bar{p}_{m-1}(\ell))\right] + \frac{\tau(\ell)}{M-1}\sum_{\tilde{m}\neq m}\left[F^H(\bar{p}_{\tilde{m}}(\ell)) - F^H(\bar{p}_{\tilde{m}-1}(\ell))\right] \\ &= [1-\tau(\ell)]\left[F^H(\bar{p}_m(\ell)) - F^H(\bar{p}_{m-1}(\ell))\right] + \frac{\tau(\ell)}{M-1}\left[1 - F^H(\bar{p}_m(\ell)) + F^H(\bar{p}_{m-1}(\ell))\right] \end{aligned} \tag{12}$$

and

$$\varphi(m, \ell) = \ell\cdot\frac{[1-\tau(\ell)]\left[F^L(\bar{p}_m(\ell)) - F^L(\bar{p}_{m-1}(\ell))\right] + \frac{\tau(\ell)}{M-1}\left[1 - F^L(\bar{p}_m(\ell)) + F^L(\bar{p}_{m-1}(\ell))\right]}{[1-\tau(\ell)]\left[F^H(\bar{p}_m(\ell)) - F^H(\bar{p}_{m-1}(\ell))\right] + \frac{\tau(\ell)}{M-1}\left[1 - F^H(\bar{p}_m(\ell)) + F^H(\bar{p}_{m-1}(\ell))\right]} \tag{13}$$

'Noise Traders'. We could imagine a third form of noise whereby a fraction of individuals receive no private signal, and therefore simply free-ride off the public information. These are analogous to the 'noise traders' that richly populate the financial literature. But they require no special treatment here, as they are subsumed in the standard model outlined in section 2. For if μ^H and μ^L have a common atom accorded the same probability under each measure, then F^H and F^L will each have an atom at 1/2. Since a noise trader is precisely someone who has the private belief equal to the common prior, namely 1/2, all results from section 4 now carry over.

Asymptotic Learning

We are now ready to investigate the effects of noise. Observe that with bounded beliefs, the interval structure of the action absorbing basins $J_1, \ldots, J_M$ in $[0, \infty)$ obtains just as before. The mere existence of action absorbing basins deserves some commentary. For one might intuit that the likelihood ratio can no longer settle down: Eventually some noisy individual will apppear and take an action so unexpected as to push the next likelihood ratio outside the putative action absorbing basin. The flaw in this logic is that precisely because the action was unexpected, the individual will be adjudged ex post to have been noisy, and his action will thus be ignored.

We now show that Theorems 3 and 5 go through unchanged.

Theorem 7 (Convergence) *Augment the standard model by one of the first two types of noise, and assume the state is H. Then $\ell_n \to \tilde{\ell}$ for some random variable $\tilde{\ell}$. If the beliefs are bounded, then $\tilde{\ell} \in J_1 \cup \cdots \cup J_M$ almost surely, while if the beliefs are unbounded, $\tilde{\ell} = 0$ almost surely.*

Proof: Since $\langle \ell_n \rangle$ is still a martingale in state H, the Martingale Convergence Theorem assures us that $\tilde{\ell}$ exists, and is almost surely finite. Let $\bar{\ell} \in \text{supp}(\tilde{\ell})$.

CASE 1: CRAZY AGENTS.

Here, the transition dynamics are given by (10) and (11), and so $\psi(m|\ell)$ is bounded away from 0, since $\tau_m > 0$ by assumption. On the other hand,

$$\varphi(m,\ell) - \ell = \ell\ \tau \frac{\left[F^L(\bar{p}_m(\ell)) - F^L(\bar{p}_{m-1}(\ell))\right] - \left[F^H(\bar{p}_m(\ell)) - F^H(\bar{p}_{m-1}(\ell))\right]}{\psi(m|\ell)},$$

and so $\varphi(m,\ell) - \ell = 0$ is satisfied under exactly the same circumstances as in the proofs of Theorems 3 and 5, because $\tau \neq 0$. Some consideration reveals that those proofs go through just as before.

CASE 2: TREMBLING AGENTS.

As with the first type of noise, all actions are taken with positive probability, and so $\psi(m|\ell)$ is indeed bounded away from 0 by (12). We wish to argue once more that $\varphi(m,\ell) - \ell = 0$ is satisfied under exactly the same conditions as in the proofs of Theorems 3 and 5. We can then use (13) to rewrite $\varphi(m,\ell) = \ell$ as follows:

$$[1-\tau(\ell)]F^L(\bar{p}_m(\ell)) + \frac{\tau(\ell)}{M-1}\left[1 - F^L(\bar{p}_m(\ell))\right] = [1-\tau(\ell)]F^H(\bar{p}_m(\ell)) + \frac{\tau(\ell)}{M-1}\left[1 - F^H(\bar{p}_m(\ell))\right]$$

which is equivalent to

$$1 - \tau(\ell) = \frac{\tau(\ell)}{M-1}$$

or simply $\tau(\ell) = 1 - 1/M$. But this violates the assumption $\tau(\ell) < 1/2$. Therefore, the proof of Theorem 3 obtains once again when $\tau(\ell) < 1/2$, while $\varphi(1,\ell) - \ell$ is bounded away on I_N, and so Theorem 5 goes through just as before. ◊

The above theorem argues that whether individuals eventually learn the true state of the world is surprisingly unaffected by a small amount of constant background noise. But the corresponding purely observational results on herding, namely Theorems 4 and 6, can no longer obtain without modification. Indeed, it is impossible that all individuals take the same action.

Let's only ask that all *rational* individuals take the same action in a herd. With such a redefinition, Theorem 4 is still valid: Herds will arise with positive probability under bounded beliefs, as ℓ must reach an action absorbing basin in finite time. But as the proof of Theorem 6 critically invokes the (now invalid) overturning principle, we cannot guarantee that a herd (correct or not) will almost surely arise.

This turns on the speed of convergence of the public belief. For a herd is tantamount to an infinite string of individuals having private beliefs that are not strong enough to counteract the public belief. Suppose that the state is H, so that we know that $q_n \to 1$ a.s. by the previous theorem. Then a correct herd arises in finite time so long as there is

not an infinite string of 'herd violators' (individuals with private beliefs below $\bar{p}_M(\ell)$). In light of the (first) Borel-Cantelli Lemma, this occurs with zero chance provided

$$\sum_{n=1}^{\infty} F^H\left(\frac{(1-q_n)\bar{r}_M}{(1-q_n)\bar{r}_M + q_n(1-\bar{r}_M)}\right) < \infty.$$

At the moment, we cannot determine whether this inequality occurs almost surely or even with positive probability. But if the public belief converges at, say, an exponential rate, then because F^H has no atom at 0 or 1 by assumption, the sum will be finite.

More States and Actions

With a denumerable action space, the only subtlety that arises is with the trembling formulation, where we shall insist upon a finite support of the tremble from any ℓ, with all those destination actions equilikely.

With more than two states the arguments go through virtually unchanged.

6. MULTIPLE INDIVIDUAL TYPES

6.1 Introduction and Motivation

Parallel to the Experimentation Literature

The results so far bear some similarity to the stylized predictions of the single-person learning theory, but are analytically much simpler. For inasmuch as individuals may ignore any future ramifications of their actions, the resulting decision problem they solve is trivial by comparison.[21] And while it is the value of information that sustains individual experimentation, observational learning by contrast is bolstered solely by the diversity of signals that subsequent individuals may entertain. There is therefore no need for ad hoc and involved methods of control theory that has dogged the experimentation literature, and greatly restricted its applicability.

Recall first an early result of this genre due to Rothschild (1974). He considered a classic economic illustration of the probabilist's 'two-armed bandit': An infinite-lived impatient monopolist optimally experiments with two possible prices each period. Rothschild showed that the monopolist would (*i*) eventually settle down on one of the prices almost surely, and (*ii*) with positive probability settle down on the less profitable price. We wish to draw some parallels with our bounded support beliefs case: The first result above corresponds to the belief cascades of Theorem 3 — for in both learning paradigms, the likelihood ratio enters an action absorbing basin, after which future signals are ignored. The second more striking pathological result corresponds to the possibility of misguided herds, as described in Theorem 4: simply put, there is always one action absorbing basin that leads individuals to adopt the most unprofitable action.

[21]Even more difficult is the marriage of the observational and experimental paradigms. For instance, while Smith (1991) explicitly tried to avoid this problem, Bolton and Harris (1993) have recently blended the two paradigms, to investigate the interplay between these forms of learning. Most notable among their findings is that when long-lived individuals' experimentation is publicly observable, there is an additional dynamic incentive to experiment, namely the desire to 'encourage' future experimentation by others.

This analogy is not without independent interest, as it foreshadows our next principal finding. For with heterogeneous preferences, an interesting new twist is introduced. As with the noise formulation, we assume that an individual's type is his private information. Yet everyone will be able to extract information from history by comparing the proportion of individuals choosing each action with the known frequencies of preference types. So long as all types do not have the same frequency, this inference intuitively ought to be fruitful. Surprisingly, however, the learning dynamics may in fact converge upon a wholly uninformative outcome, in which each action is taken with the same probability in all states. We shall argue that this 'twin pathology', which we dub *confounded learning*, arises even with unbounded private beliefs — that is, even when herding cannot occur. The essential requirement for confounded learning is that there be at least as many states of the world as actions. For otherwise, there will generically always be an action which is not taken with the same probability in each state.

Barring confounded learning, a 'herd' may arise: By this, we now mean that everyone of the same preference type will take the same action. Provided some types' vNM preferences are not identical, the overturning argument will (sometimes) fail here just as it did with noise: Unexpected actions need not radically affect beliefs, because the successors will also entertain the hypothesis that the individual was simply of a different type.

A Simple Example of Confounded Learning

There are several issues we wish to investigate, and thus find it most convenient to just consider the simplest possible specification of this model.

Assume that there are $M = 2$ actions and two preference types, labelled A and B. Individuals are of type A with chance τ_A, where the preferences of A are just as before; namely, in state H action a_2 is preferred over a_1, and conversely so in state L; type A individuals will be indifferent when their private belief equals $\bar{p}_1^A(\ell)$. Type B individuals have the opposite preferences, preferring a_1 to a_2 in state H, and conversely in state L, and having the private belief threshold $\bar{p}_1^B(\ell)$. If we assume WLOG that the state is H, then the dynamics are described by

$$\psi(1|\ell) = \tau_A F^H(\bar{p}_1^A(\ell)) + (1-\tau_A)\left[1 - F^H(\bar{p}_1^B(\ell))\right] \tag{14a}$$

$$\psi(2|\ell) = \tau_A\left[1 - F^H(\bar{p}_1^A(\ell))\right] + (1-\tau_A) F^H(\bar{p}_1^B(\ell)) \tag{14b}$$

and

$$\varphi(1,\ell) = \ell\frac{\tau_A F^L(\bar{p}_1^A(\ell)) + (1-\tau_A)\left[1 - F^L(\bar{p}_1^B(\ell))\right]}{\tau_A F^H(\bar{p}_1^A(\ell)) + (1-\tau_A)\left[1 - F^H(\bar{p}_1^B(\ell))\right]} \tag{15a}$$

$$\varphi(2,\ell) = \ell\frac{\tau_A\left[1 - F^L(\bar{p}_1^A(\ell))\right] + (1-\tau_A) F^L(\bar{p}_1^B(\ell))}{\tau_A\left[1 - F^H(\bar{p}_1^A(\ell))\right] + (1-\tau_A) F^H(\bar{p}_1^B(\ell))} \tag{15b}$$

For now, let us sidestep belief cascades as the source of incomplete learning, and simply assume that private beliefs are unbounded. Together with (14a) and (14b), this implies that $\psi(1|\ell)$ and $\psi(2|\ell)$ are bounded away from 0 for all ℓ. It is also elementary to verify that given (15a) and (15b), the two stationarity conditions $\varphi(1,\ell) = \ell$ and $\varphi(2,\ell) = \ell$

reduce to one and the same requirement: that one action (and thus the other) is be taken with equal chance in each of the two states, or

$$\tau_A F^L(\bar{p}_1^A(\ell)) + (1-\tau_A)\left[1 - F^L(\bar{p}_1^B(\ell))\right] = \tau_A F^H(\bar{p}_1^A(\ell)) + (1-\tau_A)\left[1 - F^H(\bar{p}_1^B(\ell))\right] \quad (16)$$

Intuitively, this asserts that no action reveals anything about the true state of the world, and therefore individuals simply ignore history: the private belief thresholds of the two types are precisely balanced so as to prevent successive individuals from inferring anything from history. We shall say that a solution $\hat{\ell}$ to (16) is a *confounded learning outcome*, if the interiority condition $F^H(\bar{p}_1^A(\hat{\ell})) \in (0,1)$ additionally holds.[22] This simply excludes the degenerate non-interior solutions to (16) in which the beliefs are so strong and perverse that both types are wholly convinced of their beliefs. Since a confounded learning outcome will almost surely not arise in finite time, we shall say that *confounded learning* obtains if $\ell_n \to \hat{\ell}$, where $\hat{\ell}$ is a confounded learning outcome.

Observe the following nice distinction between a belief cascade and confounded learning. In a cascade, individuals disregard their own private information and are wholly guided by the public belief. Conversely, with confounded learning, individuals (in the limit) disregard the public information and rely solely on their private signal. But while cascades and confounded learning really are different phenomena, both are pathological outcomes: Social learning stops short of an arbitrarily focused belief on the true state of the world. In the first case, the termination is abrupt, while in the second, learning slowly dies out.

Have we catalogued all possible learning pathologies? One might also imagine an altogether different conclusion of the learning dynamics. Indeed, since we have a difference equation with an interior stable point, there might perchance also exist a stable cycle, i.e. a finite set of at least two non-stable points such that the process once in the cycle would stay within the cycle. But we know that such a *stochastic steady state* cannot possibly occur because the likelihood ratio is known to converge: That $\langle \ell_n \rangle$ is a martingale in addition to a markov process is truly a useful property!

The Experimentation Literature Revisited

We now return to our earlier analogy to the experimentation literature. An interesting sequel to Rothschild (1974) was McLennan (1984), who permitted the monopolist the flexibility to charge one of a continuum of prices; he assumed for definiteness that the demand curve was one of two linear possibilities, either $q = a + bp$ or $q = A + Bp$.[23] To avoid trivialities, he assumed that neither curve dominated the other, i.e. they crossed at some interior and feasible pair $(\hat{p}, \hat{q})$. He showed that under certain conditions, the optimal price may well converge to $\hat{p}$, at which point, no further learning occurs. Intuitively, this corresponds to confounded learning in the observational learning model. The likelihood ratio is tending to an isolated stationary point outside the action absorbing basins. Furthermore, it could only arise because the action space was continuous, and thus the level of experimentation (in the sense of charging an informative price) could slowly peter out.

[22] It is easily verified that when ℓ is a confounded learning outcome if $F^H(\bar{p}_1^A(\ell)) \in (0,1)$ then $F^H(\bar{p}_1^B(\ell)) \in (0,1)$ likewise.

[23] Here, p is the price and q is the probability that this period's consumer buys, $b, B < 0$.

We shall comment on the possibility of using some of McLennan's insights in the discussion and example later in this section. But by our earlier remarks, we may adduce one implication already for single person learning theory. Since the likelihood ratio clearly must also constitute a conditional martingale in that paradigm too. McLennan's paper captured all possible pathological outcomes of a pure experimentation model.

6.2 Towards a Theory

We are now confronted with some key questions:

1. Must confounded learning outcomes exist in our model? Are they unique?
2. Even if a confounded learning outcome $\hat{\ell}$ exists, does confounded learning actually obtain (i.e. $\ell_n \to \hat{\ell}$)?
3. With unbounded beliefs, is there still a positive probability of complete learning, i.e. $\ell_n \to 0$ in state H and $\ell_n \to \infty$ in state L? If so, do correct herds arise?

For now, we can only offer partial answers to these questions. For the search for answers shall carry us into relatively uncharted territory on the local and global stability of stochastic difference equations.

Theorem 8 (Confounded Learning) *Assume there is more than one preference type.* (1) *Suppose that there are at least as many states of the world as actions, and that F^H and F^L are not entirely discrete distributions. Then confounded learning outcomes generically may exist, and when they do, confounded learning obtains with positive probability, and complete learning with chance less than 1. Yet with unbounded beliefs incorrect herds almost surely do not arise.*
(2) *If there are more actions than states, or if F^H and F^L are discrete distributions, then generically no confounded learning outcome exists, and learning is almost surely complete.*

Proof: We focus first on the simple case of two states, two actions, and two types, for which there are no more actions than states. We shall later argue that everything we say holds with more states, actions, and types, and in particular we shall prove our claims about the number of states and actions.

Observe that (16) is equivalent to

$$\frac{F^L(\bar{p}_1^A(\ell)) - F^H(\bar{p}_1^A(\ell))}{F^L(\bar{p}_1^B(\ell)) - F^H(\bar{p}_1^B(\ell))} = \frac{1-\tau}{\tau} \tag{17}$$

Here we can see why a confounded learning outcome generically exists precisely when the distribution functions have continuous segments. Simply fix any payoff assignment (and by implication the functions $\bar{p}_1^A$ and $\bar{p}_1^B$), and fix ℓ. Then calibrate τ so that ℓ solves (17). Of course, this turns the process on its head; τ should be held fixed while ℓ is varied. But if F^H and F^L vary continuously with p at the given ℓ, and if the left side of (17) is not locally constant in ℓ (which will be shown later), then around the solution we have just computed, there will exist a τ-neighborhood in which a confounded learning outcome

exists. For the partial converse, if F^H and F^L are entirely discrete, then the left side of (17) will only assume a finite number of values, and confounded learning outcomes will not be generic.

We shall not prove generally that confounded learning obtains with positive probability, but rather shall rely upon a later proof-by-example.

With unbounded private beliefs incorrect cascades cannot occur for the same reason as outlined in Lemma 5. Similarly, incorrect herds cannot occur. The only candidates for stationary points is $\ell = 0$ (that is, when the state is H), where all individuals take the correct action, and any confounded learning outcome where there is no herd.

Finally, with $M > 2$ actions, and any number of preference types, the confounded learning outcome will solve $M - 1$ independent equations in one variable ℓ — so that nonexistence will be generic. More generally, with multiple states of the world, the number of likelihood ratios will equal the number of states minus one. Hence, Theorems 3 and 5 can fail if there are at least as many states as actions (i.e. confounded learning may arise with positive probability), while they hold if there are more actions than states. ◇

Let us now discuss what is not addressed by the above theorem.

Consider the uniqueness of the confounded learning outcome. Note that with discrete distributions confounded learning outcomes ℓ are not unique when they exist, for (as seen in example below) there will in fact be an interval around ℓ of confounded learning outcomes simply because F^H and F^L are locally constant. Whether the confounded learning outcomes are unique modulo this exemption remains to be seen.

We now touch on the issue of whether learning in complete with at least as many states as actions. As complete learning and confounded learning are mutually exclusive events, we cannot yet prove that there is even a positive probability of the former. If the confounded learning outcome is very close to the initial belief $\ell = 1$, it is not at all implausible that it would attract all the mass. Also, with more than two states of the world, the dynamics are multidimensional, and these concerns become even more difficult to address.

Finally, we have until now focused on the learning dynamics rather than on the action choices. With bounded beliefs, we actually cannot be too ambitious in our assertions. For instance, in the trivial case above where both types eventually choose the same action, any limiting likelihood ratio $\bar{\ell}$ satisfies $F^H(\bar{p}_1^A(\bar{\ell})) = 1 - F^H(\bar{p}_1^B(\bar{\ell}))$. The overturning argument will hold and so a herd must arise in finite time. But if the two types take different actions in the limit, whether a herd arises is uncertain.

The Example Revisited: Local Stability

We now show that a confounded learning outcome will "attract mass" locally (that is, nearby likelihood ratios tend there with positive probability), in our simple model with two states, two actions, and two types.

If density functions f^H and f^L exist, then $\varphi(1, \ell)$ and $\varphi(2, \ell)$ are differentiable in ℓ. It is then fairly straightforward to see that if ℓ^* satisfies the stationarity criterion (17), we get

$$\psi(1|\ell^*)\varphi_\ell(1, \ell^*) + \psi(2|\ell^*)\varphi_\ell(2, \ell^*) = 1.$$

Or, in words, near the confounded learning outcome the expected next ℓ is approximately the current ℓ. This turns out (see Appendix C) to be the crucial ingredient for the local

stability of the confounded learning outcome.

In the specific example that we consider, $\mu^H(\sigma) = (1 - \sqrt{\sigma})/\sqrt{\sigma}$ while $\mu^L(\sigma) = 1$ for $\sigma \in (0,1)$. It is possible to then deduce quite simply that $F^H(p) = p^2$ and $F^L(p) = 2p - p^2$. Suppose that action a_2 is the default option of no investment, yielding a payoff equal to 0 with certainty, while action a_1 is an investment that is profitable in state H and unprofitable in state L, with the following payoffs:

Payoff of a_1	State H	State L
Type A	u	-1
Type B	-1	v

where $u, v > 0$.

It is now easy to calculate the belief thresholds $\bar{p}_1^A(\ell) = \ell/(u+\ell)$ and $\bar{p}_1^B(\ell) = v\ell/(1+v\ell)$. We wish to show that there exists a unique confounded learning outcome $\hat{\ell}$ that is locally stable in a certain weak stochastic sense. We shall apply Proposition C.1 of Appendix C to prove the stability claim. Observe that

$$\frac{F^L(\bar{p}_1^A(\ell)) - F^H(\bar{p}_1^A(\ell))}{F^L(\bar{p}_1^B(\ell)) - F^H(\bar{p}_1^B(\ell))} = \frac{u(1+v\ell)^2}{v(u+\ell)^2} \equiv \eta(\ell)$$

The behavior of the function η is now critical. Let's ignore the degenerate case $uv = 1$ in which $\eta \equiv 1$; for then a confounded learning outcome only exists if $r = 1/2$, which is the trivial case in which no inference from history is ever possible. It is fairly easy to see that η is strictly monotone if $uv \neq 1$. Assume for definiteness $uv < 1$, so that at most one confounded learning outcome exists. In fact, as the range of η is $(uv, 1/uv)$, a confounded learning outcome $\hat{\ell}$ definitely exists for all r close enough to 1/2. It is now a simple algebraic exercise to confirm that $\varphi_\ell(1, \hat{\ell}) = \varphi_\ell(2, \hat{\ell}) > 0$, and so Proposition C.1 now applies: There is a neighborhood around $\hat{\ell}$ such that if ℓ_n is in that neighborhood, then there is a positive chance that $\ell_n \to \hat{\ell}$.

Global Stability?

To be sure, we would prefer a "global stability" result.[24] Because a confounded learning outcome locally attracts mass, there are non-trivial specifications under which it attracts mass globally. Even if a stationary point is locally stable we must ensure that the dynamics can actually enter the stable region from outside: This is by no means trivial, as the stochastic process might oscillate across the stable region without ever landing there. But global properties of dynamical systems are in general notoriously hard to deduce, so that one ought to expect little progress on this latter front. Progress here is intertwined with a comprehension of the complicated asymptotic properties of stochastic difference equations.

[24]To the best of our knowledge, and much to our surprise, stability theory for stochastic difference equations really does not exist, and so we are coining terms here. We call a fixed point $\bar{y}$ of a stochastic difference equation *locally stable* if $\Pr(\lim_{n\to\infty} y_n = \bar{y}) > 0$ whenever $y_0 \in \mathcal{N}_{\bar{y}}$, a small enough neighborhood about $\bar{y}$. If $\Pr(\lim_{n\to\infty} y_n = \bar{y}) > 0$ for all y_0, then $\bar{y}$ is *globally stable*. Finally, if for every neighborhood $\mathcal{N}_{\bar{y}}$ of $\bar{y}$, there is a smaller neighborhood $\mathcal{N}_{\bar{y}}^*$ of $\bar{y}$ in $\mathcal{N}_{\bar{y}}$ such that $\Pr(y_n \in \mathcal{N}_{\bar{y}}\ \forall n \mid y_0 \in \mathcal{N}_{\bar{y}}^*) > 0$, then we simply call $\bar{y}$ *stable*. One might also wish to preface each of these labels by 'almost sure' if the probabilities are 1, but we shall not have occasion to make such assertions.

Yet McLennan managed to establish global stability for certain parameter specifications in his model. In a nutshell, his basic idea was to argue that whenever ℓ_n is on one side of the confounded learning outcome, then ℓ_{n+1} must be on the same side. We unfortunately can find no reason to expect that such a wonderful monotonicity property obtains here. For the example, it would suffice to prove (which we cannot) that globally $\varphi_\ell(m,\cdot) > 0$ for $m = 1, 2$. For in light of $\varphi(m, \hat{\ell}) = \hat{\ell}$, that would imply $\varphi(m, \ell)\hat{\ell}$ for $\ell\hat{\ell}$.

Signals about Types?

We have maintained in this section the working hypothesis that types were unobservable — for with perfectly observed types the analysis of section 4 applies. But consider the middle ground. Suppose, for simplicity, that after individual n moves, subsequent individuals receive an informative binary public signal as to his type. Then the dynamics are now modified, as there will be four different possible continuations, namely one for each of the possible combinations of individual n's action and of the type signal value. A confounded learning outcome will then have to solve three independent equations rather then one, and so generically confounded learning will not arise. We anticipate that the results of section 4 will obtain here, with appropriate modifications.

7. COSTLY INFORMATION

We now reconsider the theory under the reasonable assumption that information is costly to acquire. As we shall see, it is not only important whether this means the public or the private information, but also what the exact timing of the signal acquisition is. One might imagine variations on this theme allowing for endogenous information quality, but we shall avoid such side issues. Overall, costly signals will make incomplete learning, and thus incorrect herds, more likely.

1. Costly Public Information

First, assume that no information is revealed before the signals are acquired, and assume that the private signal is free, while the public information (the observation of all predecessors' actions) costs c. Then early individuals may not buy the public information and thus will wholly follow their own signal, the action history will eventually become very informative. Sooner or later an individual will find it worthwhile to buy the information. As the public information only (weakly) improves over time, all subsequent individuals will decide likewise. We have now arrived at the old model, and so the existing theory obtains.

Next suppose that the individuals can observe their private signal before they decide whether to buy the public information[25] — which they will do whenever its option value justifies its cost. But by the previous logic, the public information can only become more focused over time, and therefore it gets still more attractive, and therefore more and more

[25] While in the previous formulation, it did not matter whether the decision to acquire information was observable, it does now (since this decision reveals something about the individual's private signal too). But this variation is inessential for the qualitative description that follows, and so we do not further nuance here.

individuals will decide to buy it along the way. For each n, everyone with private beliefs in $[\theta_n, 1-\theta_n]$ will buy the signal, and others will not, where the threshold $\theta_n \to 1$. So, with bounded private beliefs all individuals will eventually buy the public information, and a herd almost surely arises. With unbounded private beliefs, the public belief in state H will converge to 1, and so full learning obtains.

2. Costly Private Information

Now suppose that the public information is free, but the private signal is at cost. Assume first that the purchase decision occurs before the public history is observed. Of course, private signals must be sufficiently worthwhile that the first individual is willing to buy the private signal at the cost c. By similar arguments, everyone after some individual N will not purchase private signals, and therefore the public information will not improve. *A herd will begin, even with unbounded beliefs!* Interestingly enough, N is known ex ante — it is not stochastic. That fact stands in contrast to all other herding results we have seen so far.

If individuals may first observe history before buying their private signal, then such purchases will continue until the public information reaches a certain threshold. The above result obtains, only this time the herd starts at a stochastic individual.

3. Costly Private and Public Information

Finally, consider the combination where both public and private information is costly. It is not hard to extrapolate from the above analysis: Individuals will initially only buy the private signal, but not the public one. After a while, they will start to find the public information attractive, and finally the public information will become so good that they will no longer buy the private signal. Of course, timing is an issue, and so the decision to buy the public information may or may not be predicated on the content of the private signal. This story of how the information is generated by the first agents, while later agents are free-riding on the public knowledge, captures an essential characteristic of the herding phenomenon.

A. CONSEQUENCES OF BAYES UPDATING

We originally derived all results in this Appendix by proofs other than the presented ones.[26] We consider the proofs offered here easier and more direct. Inspiration for the current formulation was found in the exercises to Chapter 8 in DeGroot (1970).

The setup is taken from section 2.1, except that now we assume that the prior chance that the state is H is $\xi \in (0,1)$. Hence, the private belief $p \in (0,1)$ satisfies

$$p(\sigma) = \frac{\xi g(\sigma)}{\xi g(\sigma) + (1-\xi)}.$$

Lemma A.1 *F^H and F^L have the same support.*

[26] One may derive our lemmas through the Law of Iterated Expectations, which for us says that applying Bayes' rule more than once is redundant. That gives rise to a quite strong characterization of the Radon-Nikodym derivative of F^H w.r.t. F^L, but we omit the details.

Proof: Note that $p(\sigma)$ is in the support of F^H (resp. F^L) exactly when σ is in the support of μ^H (resp. μ^L). But μ^H and μ^L are mutually a.c. and thus have the same support. ◇

Lemma A.2 *The function $F^L(p) - F^H(p)$ is weakly increasing for $p \le \xi$ and weakly decreasing for $p \ge \xi$. Moreover, $F^L(p) > F^H(p)$ except when $F^L(p) = F^H(p) = 0$ or $F^L(p) = F^H(p) = 1$.*

Proof: Observe that for all σ in the support of μ^H and μ^L we have

$$p(\sigma) < \xi \iff g(\sigma) < \xi g(\sigma) + (1-\xi) \iff g(\sigma) < 1$$

Now, when if Radon-Nikodym derivative $g = d\mu^H/d\mu^L < 1$ on a set of signals, then that set is accorded smaller probability mass under F^H than F^L. So F^L grows strictly faster than F^H on $(0, \xi)$ because its derivative is larger when it exists, and because it has larger atoms; similarly, F^L grows strictly slower than F^H on $(\xi, 1)$. Finally, in order that the above strict assertions obtain, $F^L > F^H$ necessarily in the interior of supp(F). ◇

Lemma A.3 *Assume $\xi = 1/2$. For any $p \in (0,1)$, we have the inequality*

$$F^H(p) \le \frac{p}{1-p} F^L(p)$$

Proof: First observe that for any $p \in (0,1)$ we have

$$p(\sigma) \le p \iff \frac{g(\sigma)}{g(\sigma)+1} \le p \iff g(\sigma) \le \frac{p}{1-p}$$

Simple integration then yields the desired

$$F^H(p) = \int_{p(\sigma)\le p} g(\sigma)\, d\mu^L(\sigma) \le \frac{p}{1-p} F^L(p)$$

◇

B. OVERTURNING

This appendix is devoted to the proof of Lemma 7.

Lemma B.1 (The Overturning Principle) *For any history h, if an individual optimally takes action a_m, then the updated likelihood ratio must satisfy*

$$\ell(h, a_m) \in \left[\frac{1-\bar{r}_m}{\bar{r}_m}, \frac{1-\bar{r}_{m-1}}{\bar{r}_{m-1}}\right) \tag{18}$$

Proof: Let the history h be given. Individual n uses the likelihood ratio $\ell(h)$ and his private signal σ_n to form his posterior odds $\ell(h)/g(\sigma_n)$ for state L. Since he optimally chose action a_m, we know from the definition of the $\bar{r}_m$'s and the tie-breaking rule, that

$$\bar{r}_{m-1} < \frac{1}{1+\ell(h)/g(\sigma_n)} \le \bar{r}_m$$

or, rewritten,

$$\frac{1-\bar{r}_{m-1}}{\bar{r}_{m-1}} > \ell(h)/g(\sigma_n) \geq \frac{1-\bar{r}_m}{\bar{r}_m} \tag{19}$$

Denote by $\Sigma(h)$ the set of signals σ_n that satisfy (19) for a given h.

To form $\ell(h, a_m)$, individual $n+1$ must calculate the probability that individual n would take action a_m, conditional on h. Individual $n+1$ knows that his predecessor would take action a_m exactly when his $\sigma_n \in \Sigma(h)$. So he can calculate

$$\alpha_m^H(h) = \int_{\Sigma(h)} g\, d\mu^L \tag{20a}$$

$$\alpha_m^L(h) = \int_{\Sigma(h)} d\mu^L \tag{20b}$$

and form the likelihood ratio

$$\ell(h, a_m) \;=\; \ell(h)\frac{\alpha_m^L(h)}{\alpha_m^H(h)} \tag{21}$$

But by definition of $\Sigma(h)$, we also know that[27]

$$\int_{\Sigma(h)} g\, d\mu^L > \frac{\bar{r}_{m-1}}{1-\bar{r}_{m-1}}\, \ell(h) \int_{\Sigma(h)} d\mu^L$$

$$\int_{\Sigma(h)} g\, d\mu^L \leq \frac{\bar{r}_m}{1-\bar{r}_m}\, \ell(h) \int_{\Sigma(h)} d\mu^L$$

Finally, (20a), (20b), (21), and the above imply that

$$\frac{1-\bar{r}_{m-1}}{\bar{r}_{m-1}} > \ell(h, a_m) \geq \frac{1-\bar{r}_m}{\bar{r}_m}$$

just as claimed. ◊

C. STABILITY OF A STOCHASTIC DIFFERENCE EQUATION

In this appendix, we first develop a *global stability* criterion for *linear* difference equations. We then use that result to derive a *stability* criterion for *linear* dynamics. Finally, we derive a result on *local stability* of a *nonlinear* dynamical system.

Consider linear stochastic difference equations of the following form. Let an i.i.d. stochastic process $\langle y_n \rangle$ be given, such that $\Pr(y_n = 1) = p = 1 - \Pr(y_n = 0)$. Define the stochastic process ℓ_n on $\mathbb{R}$ as follows: ℓ_0 is given, and

$$\ell_n = \begin{cases} a\ell_{n-1} & \text{if } y_n = 1 \\ b\ell_{n-1} & \text{if } y_n = 0 \end{cases} \tag{22}$$

[27] Notice that the strict inequality only survives the integration under the assumption that $\Sigma(h)$ has a positive measure under μ^L. Otherwise, both sides equal zero.

where a and b are fixed real constants. If $Y_n = \sum_{k=1}^{n} y_n$, then the solution to the difference equation is described by

$$\ell_n = a^{Y_n} b^{n-Y_n} \ell_0 \tag{23}$$

The following result is a rather straightforward generalization of the standard stability criterion for linear difference equations:

Lemma C.1 (Global Stability) *If* $|a|^p|b|^{1-p} < 1$ *then* $\ell_n \to 0$ *almost surely, while if* $|a|^p|b|^{1-p} > 1$ *then* $|\ell_n| \to \infty$ *almost surely.*

Proof: The essence of the proof lies in the fact that by the Strong Law of Large Numbers,

$$p \overset{\text{a.s.}}{=} \lim_{n\to\infty} \frac{Y_n}{n}$$

But if $|a|^p|b|^{1-p} < 1$ and $Y_n/n \to p$, then there exists $\varepsilon > 0$ and some N such that for all $n > N$,

$$|a|^{\frac{Y_n}{n}} |b|^{\frac{n-Y_n}{n}} < 1 - \varepsilon$$

Now, use (23) to see that

$$|\ell_n| = |a^{Y_n} b^{n-Y_n} \ell_0| = \left(|a|^{\frac{Y_n}{n}} |b|^{\frac{n-Y_n}{n}}\right)^n |\ell_0|$$

which in turn implies that $\ell_n \to 0$ a.s. The rest of the lemma follows similarly. $\diamondsuit$

This criterion deserve a few comments. One might imagine that the arithmetic mean, and not the geometric mean, of the coefficients, namely $pa + (1-p)b$, would determine the behavior of a *linear* system. In the standard theory of difference equations $p = 1$, and so these two averages coincide. If we reformulate the criterion by first taking logarithms, as

$$p\log(|a|)) + (1-p)\log(|b|)) < 0,$$

then this is reminiscent of stability results from the theory differential equations, and it is common for the logarithm to enter when translating from difference to differential equations.

It is straightforward to generalize Lemma C.1 to the case of more than two continuations, i.e. where y_n has arbitrary finite support. The analysis for multidimensional ℓ_n is also of some interest, but unfortunately in that case only one half of Lemma C.1 goes through. Indeed, let $\ell_n \in \mathbb{R}^n$ and assume

$$\ell_n = \begin{cases} A\ell_{n-1} & \text{if } y_n = 1 \\ B\ell_{n-1} & \text{if } y_n = 0 \end{cases}$$

where A and B are given real $n \times n$ matrices. Let $||A||$ and $||B||$ denote the operator norms of the matrices.[28] Then the following half of Lemma C.1 goes through, with nearly unchanged proof: If $||A||^p||B||^{1-p} < 1$ then $\ell_n \to 0$ a.s. Since this is the only part of Lemma C.1 that is used in the sequel, our local stability assertions will also be valid in a multidimensional space, too.

Next we shall provide a different stability result, which is helpful for the later characterization of non-linear systems. We consider system (22) once more.

[28] That is, $||A|| = \sup_{|x|=1} |Ax|$.

Lemma C.2 *If $|a|^p|b|^{1-p} < 1$ and $\mathcal{N}_0$ is any open neighborhood about 0, then there is a positive probability that $y_n \in \mathcal{N}_0$ for all n, provided $y_0 \in \mathcal{N}_0$.*

Proof: First, if a or b is 0, then there is a positive probability of jumping to 0 immediately, and the system will stay there. So, assume now that $a, b, \ell_0 \neq 0$. We already know from the previous lemma that $\ell_n \to 0$ almost surely. To be explicit, this means that for all $N \in \mathbb{N}$, we have

$$\Pr\left(\bigcup_{M\in\mathbb{N}} \bigcap_{n>M} \{\omega \in \Omega^H : ||\ell_n|| < 1/N\}\right) = 1$$

There must then be some $M \in \mathbb{N}$ such that

$$\Pr\left(\{\omega \in \Omega^H : \forall n \geq M, ||\ell_n|| < 1/N\}\right) > 0$$

In particular, so long as $\ell_M \in (-1/N, 1/N)$, there is a positive chance of $\ell_n \in (-1/N, 1/N)$ for all $n \geq M$. Also, for given ℓ_0, $\text{supp}(\ell_M)$ is finite, and $\ell_0 \neq 0$ implies $\ell_M \neq 0$; therefore, there exists $\bar{\ell}_M \in (-1/N, 1/N)\setminus\{0\}$ such that $\ell_n \in (-1/N, 1/N)$ for all $n \geq M$ if $\ell_M = \bar{\ell}_M$.

The proof is closed by appealing to two simple invariance properties of the problem. First, time invariance allows us to conclude that $\ell_n \in (-1/N, 1/N)$ for all n if $\ell_0 = \bar{\ell}_M$. Second, the equations are linear, so that if $\ell_0 < \bar{\ell}_M$ then $\ell_n \in (-(\ell_0/\ell_M)/N, (\ell_0/\ell_M)/N)$ for all n. But finally notice that from any point of a large neighborhood, one can reach the inner neighborhood in only a finite number of steps, which occurs with positive probability too. This completes the proof. ◇

We next use these results to investigate the local stability of non-linear stochastic dynamical systems. While we have so far has proceeded on a very general level, we now take advantage of the far more special assumptions relevant for the application in section 6. This is only slight overkill, since the appendix does not require the martingale property.

In the notation of section 3, the system is now described by

$$\ell_n = \begin{cases} \varphi(1, \ell_{n-1}) & \text{if } y_n = 1 \\ \varphi(2, \ell_{n-1}) & \text{if } y_n = 0 \end{cases} \tag{24}$$

where $\varphi(1, \cdot)$ and $\varphi(2, \cdot)$ are given functions, and ℓ_0 is given. Moreover, we shall now abandon the i.i.d. assumption on the stochastic process $\langle y_n \rangle$, and posit instead that

$$\Pr(y_n = 1) = \psi(1|\ell_{n-1}) = 1 - \Pr(y_n = 0)$$

We are concerned with stationary points $\hat{\ell}$ of (24), namely the solutions to

$$\varphi(1, \hat{\ell}) = \hat{\ell} \text{ and } \varphi(2, \hat{\ell}) = \hat{\ell} \tag{25}$$

Proposition C.1 (Local Stability) *Fix a stationary point $\hat{\ell}$ given by (25). Assume that at $\hat{\ell}$, $\varphi(1, \cdot)$ and $\varphi(2, \cdot)$ are both continuously differentiable, while $\psi(1|\cdot)$ is continuous. Suppose also that $\hat{\ell}$ satisfies $\psi(1|\hat{\ell}) \in (0, 1)$, $\varphi_\ell(1, \hat{\ell}) > 0, \varphi_\ell(2, \hat{\ell}) > 0$, $\varphi(1, \hat{\ell}) \neq \varphi(2, \hat{\ell})$, and $\psi(1|\hat{\ell})\varphi_\ell(1, \hat{\ell}) + (1 - \psi(1|\hat{\ell}))\varphi_\ell(2, \hat{\ell}) = 1$. Then there is a neighborhood around $\hat{\ell}$, such that from any point in this neighborhood there is a positive probability that $\ell_n \to \hat{\ell}$.*

Proof: We proceed as follows. First, we linearize (and majorize) the nonlinear dynamical system around $\hat{\ell}$ by a linear stochastic difference equation of the form just treated (that satisfies the conclusions of Lemma C.2). Next we argue that the conclusion of Lemma C.2 must apply to the original non-linear dynamical system.

Since $\psi(1|\hat{\ell})\varphi_\ell(1,\hat{\ell}) + (1-\psi(1|\hat{\ell}))\varphi_\ell(2,\hat{\ell}) = 1$ and $\varphi_\ell(1,\hat{\ell}) \neq \varphi_\ell(2,\hat{\ell})$, the arithmetic mean-geometric mean inequality yields

$$\varphi_\ell(1,\hat{\ell})^{\psi(1|\hat{\ell})}\varphi_\ell(2,\hat{\ell})^{(1-\psi(1|\hat{\ell}))} < 1. \tag{26}$$

By the continuity assumptions on $\varphi_\ell(1,\cdot)$ and $\varphi_\ell(2,\cdot)$, this inequality obtains in a neighborhood of $\hat{\ell}$. As $\varphi_\ell(1,\hat{\ell}) \neq \varphi_\ell(2,\hat{\ell})$ and their average is 1, we may assume WLOG that $\varphi_\ell(1,\hat{\ell}) < 1$.

We now claim that there are constants $a, b, p > 0$, and a small enough neighborhood $\mathcal{N}(\hat{\ell})$ around $\hat{\ell}$, such that for all $\ell \in \mathcal{N}(\hat{\ell})$:

$$\begin{gathered} a^p b^{1-p} < 1 \\ 0 < \varphi_\ell(1,\ell) < a < 1 < \varphi_\ell(2,\ell) < b \\ 0 < p < \psi(1|\ell) \end{gathered}$$

This may not be obvious, so let us spell it out. Over any compact neighborhood of $\hat{\ell}$ we can separately maximize $\varphi_\ell(1,\ell), \varphi_\ell(2,\ell)$ and $-\psi(1|\ell)$. If we substitute the three maxima a, b, and $-p$ respectively, in (26), then continuity tells us that its left side converges to

$$(\varphi_\ell(1,\cdot))^p(\varphi_\ell(2,\cdot))^{(1-p)}$$

as the compact neighborhoods get small enough.

Fix $\ell_0 \in \mathcal{N}(\hat{\ell})$. There clearly exists a sequence of i.i.d. stochastic variables $\langle\sigma_n\rangle$, each uniformly distributed on $[0,1]$, such that $y_n = 1$ exactly when $\sigma_n \leq \psi(1|\ell_{n-1})$ and $\tilde{y}_n = 0$ otherwise. Introduce a new stochastic process $\langle\tilde{y}_n\rangle$ defined by $\tilde{y}_n = 1$ when $\sigma_n \leq p$, and $\tilde{y}_n = 0$ otherwise. Use this to define a new stochastic process $\langle\tilde{\ell}_n\rangle$ by $\tilde{\ell}_0 = \ell_0 - \hat{\ell}$, and

$$\tilde{\ell}_n - \hat{\ell} = \begin{cases} a(\tilde{\ell}_{n-1} - \hat{\ell}) & \text{if } \tilde{y}_n = 1 \\ b(\tilde{\ell}_{n-1} - \hat{\ell}) & \text{if } \tilde{y}_n = 0 \end{cases}$$

We now argue that the linear process $\langle\tilde{\ell}_n\rangle$ majorizes the non-linear system $\langle\ell_n\rangle$. Observe that $\langle\tilde{y}_n\rangle$ are independent because $\langle\sigma_n\rangle$ are. Thus Lemma C.1 is valid, and tells us that $\tilde{\ell}_n \to \hat{\ell}$ a.s., while Lemma C.2 asserts that there is a positive probability that $\tilde{\ell}_n \in \mathcal{N}(\hat{\ell})$ for all n. So consider just such a realization of $\langle\sigma_n\rangle$ whereby $\tilde{\ell}_n \in \mathcal{N}(\hat{\ell})$ for all n and $\tilde{\ell}_n \to 0$. Because $\psi(1|\ell) > p$ when $\ell \in \mathcal{N}(\hat{\ell})$, we have $\tilde{y}_n = 1 \Rightarrow y_n = 1$. Thus $0 < \varphi_\ell(1,\ell) < a < 1 < \varphi_\ell(2,\ell) < b$ yields the desired majorization for all n (in this realization), namely

$$||\tilde{\ell}_n - \hat{\ell}|| \geq ||\ell_n - \hat{\ell}||$$

For no matter what is the outcome of y_n, $\tilde{\ell}_n$ is always moved further away from $\hat{\ell}$ than is ℓ_n (as seen in figures 2 and 3). So, for any such realization of $\langle\sigma_n\rangle$, $\ell_n \to \hat{\ell}$. We thus conclude that $\ell_n \to \hat{\ell}$ with positive probability. $\diamondsuit$

Figure 2: **Dominance Argument I.** This depicts how an iteration under $\varphi(1,\cdot)$ brings the image closer to the stationary point $\hat{\ell}$ than an iteration under the linearization with slope a, when $\varphi_\ell(1,\cdot) < a$.

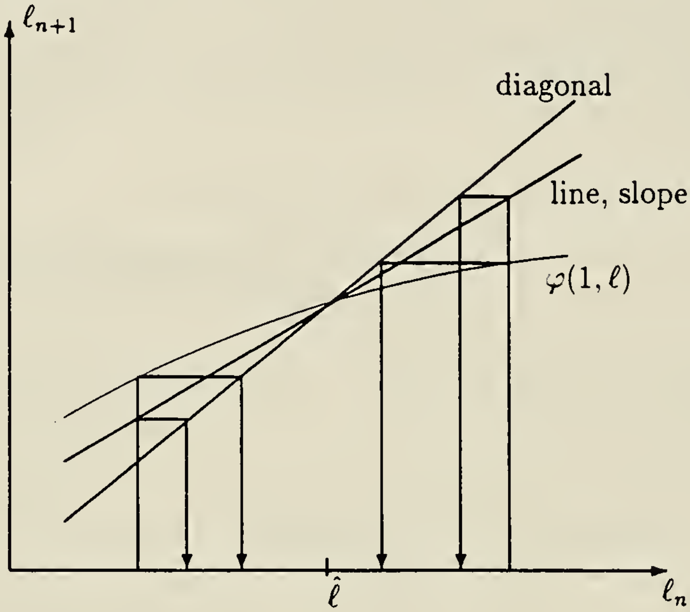

Figure 3: **Dominance Argument II.** This depicts how an iteration under $\varphi(2,\cdot)$ moves the image point closer to the stationary point $\hat{\ell}$ than an iteration under the linearization with slope b, when $\varphi_\ell(2,\cdot) < b$.

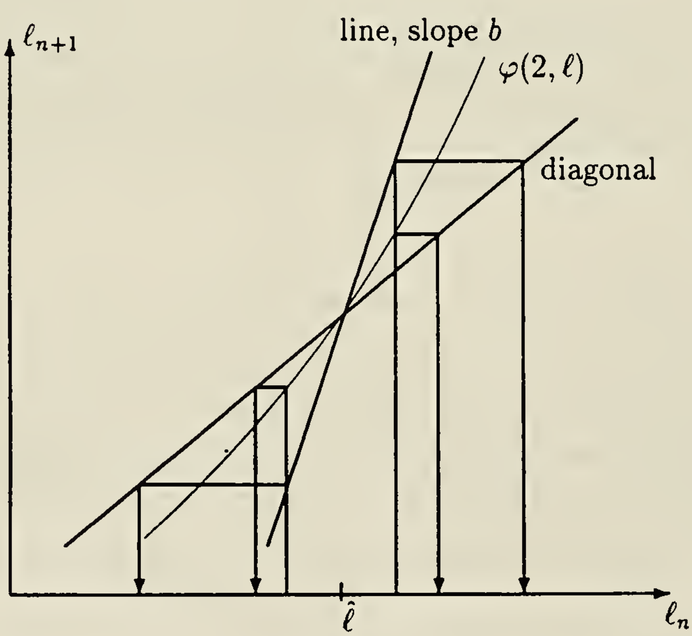

References

Aghion, Philippe, Patrick Bolton, Christopher Harris, and Bruno Jullien, "Optimal Learning by Experimentation," *Review of Economic Studies*, 1991, *58* (196), 621–654.

Arthur, W. Brian, Yu. M. Ermoliev, and Yu. M. Kaniovski, "Strong Laws For a Class of Path-Dependent Stochastic Processes with Applications," in I. Arkin, A. Shiraev, and R. Wets, eds., *Stochastic Optimization : Proceedings of the International Conference, Kiev, 1984*, New York: Spinger-Verlag, 1986.

Avery, Chris and Zemsky, "Herding in Financial Markets," 1992. PhD Thesis, Chris Avery, Stanford University.

Banerjee, Abhijit, "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 1992, *107*, 797–817.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch, "A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades," *Journal of Political Economy*, 1992, *100*, 992–1026.

Billingsley, Patrick, *Probability and Measure*, New York: Wiley, 1986.

Bolton, Patrick and Chris Harris, "Strategic Experimentation," 1993. Nuffield College mimeo.

Bray, Margaret and David Kreps, "Rational Learning and Rational Expectations," in G. Feiwel, ed., *Kenneth Arrow and the Ascent of Economic Theory*, London: Macmillan, 1987, pp. 597–625.

Breiman, Leo, *Probability*, Reading, Mass.: Addison-Wesley, 1968.

Chamley, Christophe and Douglas Gale, "Information Revelation and Strategic Delay in a Model of Investment," 1992. Boston University mimeo.

DeGroot, Morris H., *Optimal Statistical Decisions*, New York: McGraw-Hill, 1970.

Futia, Carl A., "Invariant Distributions and the Limiting Behavior of Markovian Economic Models," *Econometrica*, 1982, *50*, 377–408.

Harsanyi, J., "Games with Incomplete Information Played by Bayesian Players," *Management Science*, 1967–68, *14*, 159–182, 320–334, 486–502.

Jovanovic, Boyan, "Micro Shocks and Aggregate Risk," *Quarterly Journal of Economics*, 1987, pp. 395–409.

Lee, In Ho, "On the Convergence of Informational Cascades," *Journal of Economic Theory*, 1993, *61*, 395–411.

McLennan, Andrew, "Price Dispersion and Incomplete Learning in the Long Run," *Journal of Economic Dynamics and Control*, 1984, *7*, 331–347.

Rothschild, Michael, "A Two-Armed Bandit Theory of Market Pricing," *Journal of Economic Theory*, 1974, *9*, 185–202.

Rudin, Walter, *Real and Complex Analysis*, 3rd ed., New York: McGraw-Hill, 1987.

Selten, Reinhardt, "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, 1975, *4*, 25–55.

Smith, Lones, "Error Persistence, and Experiential versus Observational Learning," 1991. Foerder Series working paper.

– and Peter Sørensen, "An Example of Non-Martingale Learning," 1994. In progress.